


GENERIC SATELLITE MONITORING
EXPERT SYSTEM

THESIS

Loretta A. Kelemen, Captain, USAF

AFIT/GSO/ENG/94D-02



DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

19941228 015

AFIT/GSO/ENG/94D-02

GENERIC SATELLITE MONITORING
EXPERT SYSTEM

THESIS

Loretta A. Kelemen, Captain, USAF

AFIT/GSO/ENG/94D-02

INTO GREAT BRITAIN

Approved for public release; distribution unlimited

GENERIC SATELLITE MONITORING EXPERT SYSTEM

THESIS

Presented to the Faculty of the Graduate School of Engineering

of the Air Force Institute of Technology

Air University

In Partial Fulfillment of the

Requirements for the Degree of

Master of Science in Space Operations

Loretta A. Kelemen, B.S.

Captain, USAF

December, 1994

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

Approved for public release; distribution unlimited

Acknowledgments

This research was a true challenge and would not have been successful without the help of certain individuals. Foremost, I would like to thank my advisor, Major Gregg Gunsch, who provided continuous support and direction throughout this effort. He consistently pointed me in the proper direction when I would stray, and his added deadlines were very helpful in ensuring completion of the thesis and I would recommend such a schedule along with practice Defenses for future graduates. I would also like to thank my committee members, Col. Gregory Parnell and Dr. Eugene Santos for their contributions to this effort.

I believe the support of an involved sponsor adds greatly to the success of a thesis, and my sponsor, Lt Col Crowley of Phillips Laboratory, was very involved and helpful throughout this research and I would like to thank her for her support.

Of course, the thesis would never be completed without the computer support of Mr. Dan Zambon and Captain Gary Mauersberger. These two gentlemen were a tremendous help in maintaining my computer at peak performance and many thanks go to them.

To my school friends who listened tentatively to my continuous complaints and sorrows, I thank you for your ear. And to my best friend of all, my husband Vince, thanks for keeping everything in perspective.

Most of all, I thank God for the infinite number of blessings he has given me, whom without, nothing is possible.

Loretta A. Kelemen

Table of Contents

	Page
Acknowledgments	ii
List of Figures	vii
List of Tables	viii
Abstract	ix
I. Introduction	1-1
1.1 Background	1-1
1.2 Current Satellite Operation Issues	1-5
1.2.1 Satellite Operator Competence	1-5
1.2.2 Information Overload	1-5
1.2.3 Blue Suit Syndrome.....	1-5
1.2.4 Non-Standardization.....	1-6
1.3 Motivation	1-6
1.4 Research Perspective.....	1-7
1.5 Problem	1-8
1.6 Research Objectives	1-10
1.7 Approach.....	1-11
1.8 Scope.....	1-11
1.9 Executive Overview	1-13
II. Literature Review	2-1
2.1 Background	2-1
2.2 Research and Development Programs in Intelligent Satellite Control	2-2
2.2.1 StarPlan/Paragon	2-2
2.2.1.1 StarPlan I.....	2-3
2.2.1.2 Paragon.....	2-4
2.2.1.3 StarPlan II.....	2-6
2.2.1.4 StarPlan III	2-6

	Page
2.2.2 MARVEL.....	2-7
2.2.3 SHARP	2-7
2.2.4 ASW	2-8
2.2.5 MARPLE	2-9
2.2.6 SAGE and Rule Reuse.....	2-10
2.2.6.1 Architecture	2-10
2.2.6.2 Rule Reuse.....	2-12
2.3 AI Reasoning Methodologies	2-13
2.3.1 Case-Based Reasoning.....	2-14
2.3.2 Rule-Based Reasoning	2-14
2.3.3 Model-Based Reasoning	2-17
2.3.4 Blending Different Reasoning Methodologies	2-18
2.3.5 Summary	2-19
2.4 Conclusion.....	2-19
III. Methodology	3-1
3.1 Overview	3-1
3.2 Domain.....	3-2
3.2.1 The Telemetry, Tracking, and Commanding Subsystem	3-3
3.2.1.1 Comparison of the TT&C Uplink Sections	3-3
3.2.1.2 Comparison of the TT&C Downlink Sections.....	3-5
3.3 Development Environment.....	3-6
3.4 Extracting Commonalities	3-7
3.5 Rule-Base Design Considerations.....	3-8
3.6 Summary.....	3-10
IV. Design and Implementation	4-1
4.1 Overview	4-1
4.2 MAGIC, TAS, and GISMO	4-1
4.3 Work-Arounds.....	4-2
4.4 Supporting Databases.....	4-5
4.4.1 The Telemetry Database	4-5
4.4.2 The Limits/Status Database	4-5
4.4.3 The Ephemeris Database.....	4-7

	Page
4.5 Paradox Scripts.....	4-8
4.5.1 The Expected State Script	4-8
4.5.2 The Satellite State Script	4-8
4.5.3 The Update Script	4-9
4.6 Generic Rule-Base Design.....	4-10
4.6.1 Special Rule Structure Requirements	4-14
4.7 GISMO Operations.....	4-16
4.7.1 GISMO's Limitations.....	4-16
4.8 Summary.....	4-17
V. Results and Issues.....	5-1
5.1 Introduction.....	5-1
5.2 GISMO, the Initial Two-Satellite Prototype	5-1
5.2.1 Telemetry Commonalities	5-1
5.2.2 Prototype Limitations to Overcome	5-5
5.2.2.1 Multiple Point Failures	5-5
5.2.2.2 Communication Link Drop-Outs	5-5
5.2.2.3 100% Recovery Success, Hard Failures, and Immediate Satellite Response	5-6
5.2.2.4 Memory	5-7
5.2.3 Concerns and Opinions	5-10
5.2.3.1 Concerns.....	5-10
5.2.3.2 Opinions	5-11
5.2.4 Items to Investigate	5-12
5.2.4.1 More Powerful Software Tools	5-12
5.2.4.2 Telemetry Preprocessing	5-13
5.2.4.3 Additional Suggestions.....	5-15
5.3 Initial GISMO Prototype with a Third Satellite Extension.....	5-16
5.3.1 TT&C Differences	5-17
5.3.1.1 Operating Modes	5-17
5.3.1.2 Receiver Operation	5-18
5.3.1.3 Cross-Strap Timer.....	5-18
5.3.1.4 Encrypter Cross-Strap.....	5-19
5.3.2 TT&C Similarities	5-19
5.3.2.1 GPS Telemetry Commonalities.....	5-19
5.3.3 Merging Difficulties.....	5-20
5.3.4 Results and Findings	5-23

	Page
5.4 Summary.....	5-25
VI. Summary and Conclusions.....	6-1
Appendix A. The TT&C Subsystem.....	A-1
A.1 DSCS II TT&C Uplink Section	A-1
A.2 DSCS III TT&C Uplink Section.....	A-4
A.3 DSCS II TT&C Downlink Section	A-8
A.4 DSCS III TT&C Downlink Section	A-9
A.5 GPS TT&C Subsystem.....	A-10
A.5.1 GPS TT&C Uplink Section.....	A-11
A.5.2 GPS TT&C Downlink Section	A-14
Appendix B. Data Representation Schemes.....	B-1
B.1 GISMO Variable/Qualifier Format.....	B-1
B.1.1 Actual Telemetry	B-1
B.1.2 Expected Telemetry	B-1
B.1.3 Availability of Redundant Units.....	B-1
B.1.4 Qualifiers to Represent Out-of-Limit Conditions	B-1
Appendix C. Detailed Analysis.....	C-1
C.1 DSCS II and DSCS III Telemetry Comparison	C-1
C.2 GPS Telemetry Comparison	C-4
C.3 Detectable and Undetectable TT&C Anomalies	C-7
C.3.1 DSCS II Anomalies.....	C-7
C.3.2 DSCS III Anomalies	C-9
C.3.3 GPS Anomalies.....	C-11
Bibliography.....	BIB-1
Vita.....	VITA-1

List of Figures

Figure	Page
1.1 Satellite Command and Control.....	1-4
1.2 The MAGIC Architecture	1-9
1.3 The DSCS II and DSCS III satellites.....	1-12
4.1 The GISMO Prototype Architecture	4-3
4.2 DSCS II Telemetry Database Excerpt	4-5
4.3 DSCS II Limits/Status Database Excerpt.....	4-6
4.4 9446 Ephemeris Database	4-7
4.5 GISMO Generic Rule Structure	4-12
A.1 Uplink Section of the DSCS II TT&C Subsystem.....	A-2
A.2 Uplink Section of the DSCS III TT&C Subsystem.....	A-6
A.3 Downlink Section of the DSCS II TT&C Subsystem	A-8
A.4 Downlink Section of the DSCS III TT&C Subsystem.....	A-9
A.5 The GPS TT&C Subsystem.....	A-12

List of Tables

Table	Page
5.1 DSCS II and DSCS III TT&C Commonalities	5-1
5.2 Generic Rule Variables	5-3
5.3 Combined Telemetry Points.....	5-14
5.4 Commonalities After Telemetry Preprocessing	5-15
5.5 GPS Commonalities with DSCS II and DSCS III	5-20
5.6 Generic Rule Compatibility with GPS.....	5-23

Abstract

The domain of satellite operations is undergoing major changes. Satellite operators are no longer receiving detailed satellite training, instead they are taught the fundamentals necessary to command and control various multi-million dollar satellites. The need is clear: an efficient and economical automated system is necessary to assist the current satellite operator in the daily tasks of maintaining the health and status of these high priority DOD resources.

Intelligent satellite controllers have been under research and development since the early 1980s to meet this need. All of these systems, however, have focused on the command and control of one particular constellation of satellites. In a military striving for more efficiency and lower costs, the idea of developing a unique intelligent controller system for each satellite constellation is unaffordable.

This research provides support, through the development of a prototype expert system, for the concept of a generic intelligent satellite controller. This capability would allow a generic rule-base to operate and maintain multiple satellite systems.

An initial prototype was developed to detect anomalies on one subsystem of two different satellites. After the two-satellite prototype was created a third satellite was analyzed to show support for the viability of the two-satellite prototype.

More research is necessary before making the final decision on the feasibility of a generic intelligent satellite controller, but this thesis has created some support for the concept and laid the foundation for future extensions.

GENERIC SATELLITE MONITORING EXPERT SYSTEM

I. Introduction

1.1 Background

Air Force Space Command controls over eighty operational DOD satellites using a large, expensive, and hard-to-maintain support segment. It takes 4,500 personnel, working with expensive and complex computer equipment, to maintain these multi-million dollar satellites (23:500-503). In the future, the number of satellites will continue to grow while the number of support personnel will remain constant, if not decline (1:4-1).

Satellites are categorized by their particular mission, these being Navigation, Communication, Warning, or Remote Sensing. Under each mission category, the satellites are further grouped into constellations. The entire constellation is not built and launched at one time. These satellites are usually built and launched in groups of two, so it may take many years before a full constellation is in orbit. This time lag in manufacture and launch allows for new technology and previous lessons learned to be incorporated into later models. Each constellation of satellites contains one or more of these models. Therefore, it is easy to see that nearly every active satellite is different from every other in some way. These differences in missions, constellations, and models within constellations of satellites require specific, unique, operational support.

Resources are located around the world to support these active satellites. The resources include satellite ground antennas, operation centers, and communication links. Within the operation centers there are specially-trained operators. Training these specialized operators to maintain a constellation of satellites takes an average of one year, during which the operator is trained on all different models within a particular constellation.

Because of the varied complexity of satellite operations, only officers with technical degrees were trained as satellite operators in the early days of Air Force satellite operations. The officers received very detailed satellite-specific training, to include training about the satellites' circuit diagrams. Some time later, Air Force Space Command decided it could no longer afford to pay officers to be satellite operators. The officers were then tasked with putting their knowledge and experience into checklist and operation manual format. Once developed, this documentation would be used to help young airmen perform satellite operations. Air Force Space Command also reduced the required training time since most of the information necessary to perform satellite operations could be found in these new manuals.

The current training program follows a black box training methodology. The new satellite operators are trained to know what the inputs and outputs of a black box should look like, but not what takes place within the black box. The decision-makers believe knowledge of the internal workings of the black box is not necessary to perform routine satellite operations. An engineering shop was formed to handle any unusual situation not covered in the checklists or operation manuals. The engineers complete the same training as the new satellite operators followed by another training program once in the engineering shop.

If an anomaly occurs for which there is no pre-defined passplan (recovery procedure), the satellite engineer is called upon to solve the anomaly. The procedure these engineers follow first ensures the satellite is in a safe configuration. They then begin a cautious, but rigorous, process to troubleshoot the problem. Once the anomaly is diagnosed, the engineer creates a new passplan to resolve the problem. Afterwards, the passplan is published to be used again if the same anomaly reoccurs on a similar vehicle.

In October of this year, Air Force Space Command changed its satellite operational concept from the one discussed above, in which the operator is trained on one particular satellite constellation, to one in which the operator is trained on the very basic fundamentals common to multiple satellite constellations. The satellite operator will use this fundamental

training, along with detailed satellite-specific checklists and operation manuals, to operate and maintain multiple constellations of satellites. This decision is the result of current Air Force reductions. These reductions have forced Air Force Space Command to strive to do "More With Less," yet the current job performance standards must be maintained. The satellites being controlled by these operators are high priority government systems. A failure of one of these satellites could affect the lives of many military members that rely on the capabilities of these satellites in their military mission. Therefore, the operational status of each satellite must be maintained at its highest possible state throughout the satellite's operational lifetime. Different constellations are designed for different life spans, ranging from several months to ten years.

To ensure these systems operate at peak performance throughout their life, satellite operators remotely monitor the health and status of these vehicles daily. This means the satellite operator establishes a communications link from his or her operations center to a ground satellite antenna in visible range of the satellite to be monitored, as pictured in Figure 1.1. These satellite ground antennas are part of the Air Force Satellite Control Network (AFSCN). The AFSCN manages and maintains over fifteen of these ground antennas at locations around the world.

After the ground link is established, azimuth and elevation data is sent to the ground antenna to point it toward the satellite. Once pointing toward the satellite, the ground antenna will begin to collect telemetry (sensor) data being transmitted by the satellite, if the satellite transmitter is on. The telemetry includes data such as currents, voltages, and temperatures of on-board components, along with discrete status readings. At the operations center, the telemetry is preprocessed into an alphanumeric format and displayed on the satellite operator's computer consoles. The operator analyzes the telemetry to determine the satellite's state-of-health. The amount of telemetry the operator analyzes can range from two hundred to over seven hundred sensor data readings. If the satellite is in a critical situation, the satellite operator has, on average, ten minutes to determine the state-of-health of the satellite and react

to the situation. A physical separation, ranging from 200 to 20,000 nautical miles between the operator and the satellite, adds more difficulty to the diagnosing task. The operator must base his or her decision on a mental picture derived from snapshots of incomplete satellite telemetry data obtained at each satellite contact (10:1).




Figure 1.1 Satellite Command and Control

When a satellite is monitored it is commonly called a support or a contact. A contact with a satellite can range anywhere from five minutes to several hours. Some satellite constellations may require two to three supports per day while others may need continuous support. In addition to monitoring satellite state-of-health, a contact may be scheduled for one or more of the following reasons: position data collection, nominal commanding, satellite maneuver, satellite reconfiguration, and mission data collection.

1.2 Current Satellite Operation Issues

The above discussion has described the evolution of satellite operations in Air Force Space Command. The evolutionary changes in the operational concept, along with the increasing complexity of satellite systems has led to some issues that must be resolved to ensure the high quality of satellite operations is maintained. This section will describe some of the most important issues forming the basis of this research.

1.2.1 Satellite Operator Competence. Satellite operations is very tedious work and the job performance expectations are extremely high. Satellite operators must be remarkably competent in their job to perform within expected Air Force standards. These standards do not allow for mistakes since even minor human error can easily cost tens of millions of dollars in irrecoverable satellite damages. The young airman operators are not equipped to maintain this high level of job performance. We must better equip them to enhance their operational capability.

1.2.2 Information Overload. As new satellites grow in capabilities, they also grow in complexity. Some satellites are so complex a single human is incapable of managing the entire satellite. These complex satellites transmit an enormous amount of telemetry. Operators of these systems are inundated with information and can no longer effectively diagnose the satellite's state-of-health without some form of assistance.

1.2.3 Blue Suit Syndrome. All military branches of service periodically reassign their personnel. The officer corps is reassigned every three to five years, while the enlisted force is moved even more frequently. The military operates this way to give their members a broader background so they can be better leaders. Therefore, the frequent reassignment process exists to maintain efficient and energetic job performance from every military member.

On the other hand, large amounts of money are spent to train the member at each new assignment. Once trained, the member becomes proficient and gains valuable experience on his or her job. When this highly trained, experienced, and knowledgeable member is reassigned, all of this knowledge is lost and the cycle continues with more money spent to train the next member. This has become a costly concern for an Air Force faced with large financial cutbacks. The military needs a computer system to capture and maintain the expertise of these military members before they are reassigned.

1.2.4 Non-Standardization. An issue that directly contributes to the high costs of satellite operations and maintenance is non-standardization. Currently, there is no standardization between designs of different satellite constellations. This leads to non-standardization of the support segment, the ground operation centers, and the software, hardware and operators required to maintain multiple satellite constellations. The United States Space Command is very concerned with the problems and costs related to non-standardization. The Integrated Satellite Control (ISC) Standards Management Committee was formed to "Advocate the adoption or development of appropriate standards within the domain of ISC to promote interoperability across ALL NORAD/USSPACECOM Systems." (12:2)

1.3 Motivation

The Air Force realizes the expense and complexity of maintaining the current satellite support segment and wishes to find a way to improve the current way of doing business through the use of a more manageable, cost effective and maintainable system. Many of the problems discussed in this chapter could be eliminated with the implementation of a satellite expert system to automate much of the current satellite operation process.

The time has come where individual operators can no longer effectively control and maintain satellite systems without the assistance of computers. We must have intelligent assistants to aid the satellite operator in his or her tasks. These intelligent systems will contain the knowledge of the most competent experts, yet be flexible, maintainable and user friendly. The system will not only contain the knowledge of experts, but new knowledge gained from experience should be easily added. Therefore, the knowledge of the assistant would grow and there would be no worry of losing this knowledge due to a reassignment. Last, the intelligent computer could easily process the large amounts of data streaming into the operations center from these complex satellites.

There has been considerable research in the domain of satellite operations to resolve the above described problems. Researchers have built intelligent systems to automate satellite operations. These systems contain the expertise and capabilities to process large amounts of data. The one issue overlooked in these research efforts is non-standardization. All of the current research is focused on a particular satellite constellation and are not generic as to allow operation on satellites of different constellations. The need to create a generic, cost-efficient, intelligent satellite controller is the motivation for this research.

1.4 Research Perspective

This research effort is a small part of a much larger, multi-year project supported and managed by the Satellite Control and Simulation Division of Phillips Laboratory at Kirtland Air Force Base in Albuquerque, New Mexico. The project name is MAGIC, which stands for Multimission Advanced Ground Intelligent Control. The MAGIC system will be capable of managing and controlling multiple satellite constellations, easily adapting to new constellations, improving operator effectiveness, and enhancing operational capabilities. It will do this through object-oriented database techniques, knowledge reuse, and automated reasoning methods. This thesis effort investigates the

effectiveness of knowledge reuse which will play a key role in the cost effectiveness of the MAGIC project.

Recently, Phillips Laboratory was tasked by the 3rd Satellite Operations Squadron (3SOPS), located at Falcon AFB Colorado, to design a Telemetry Analysis System (TAS). The 3SOPS is in the middle of an operational concept change and will be replacing their current operators trained on specific satellites with generic trained operators. They desire the TAS system to aid the transition process by providing additional, easy-to-understand, interpretations of the satellite telemetry data. Phillips Laboratory is taking this opportunity to begin the early architecture designs of what will one day be MAGIC.

TAS is a real-time telemetry analysis system which processes raw satellite telemetry and checks for any Out-Of-Limit (OOL) conditions. It also offers real-time telemetry plots and graphs, and trending analysis for the satellite operator. No diagnostic capabilities are planned for the first cut of the TAS system, nor are any generic knowledge reuse concepts included.

This is where GISMO, the Generic Satellite Monitor, fits into the large picture. GISMO is the prototype created in this thesis effort to investigate the development of a generic reusable rule-base. The success of this research will help ease the transition of TAS into the planned MAGIC architecture. Figure 1.2 is a pictorial representation of the proposed MAGIC architecture. Section 4.2 discusses the relationship of GISMO and MAGIC.

1.5 Problem

An efficient and economical automated system is necessary to assist the current satellite operator in the daily tasks of maintaining the health and status of high priority DOD resources. A successful generic reusable rule-base could reduce the cost of developing and maintaining such an automated system.




Figure 1.2 The MAGIC Architecture

1.6 Research Objectives

This research effort tested the following hypothesis:

Main Hypothesis:

A generic satellite monitoring expert system can perform satellite anomaly detection on satellites of different constellations.

The primary purpose of this research was to determine if a generic expert system could successfully detect system malfunctions on satellites of different constellations. Such an expert system could revolutionize the way satellite operations is performed today.

To determine the validity of the primary hypothesis, two sub-hypotheses were developed to give direction to the research effort.

Commonalities Sub-hypothesis:

There exist a significant number of commonalities between satellites of different constellations.

This sub-hypothesis states that there are enough commonalities between different satellites to support development of a generic satellite expert system. The generic knowledge base was built to diagnose problems on the systems common among different satellites, while satellite-specific systems were diagnosed using a satellite-specific knowledge base. The size of the generic knowledge base must be large enough to justify its existence. The decision of what is large enough is not one that this research can provide. The purpose of this research is to provide the necessary data to the sponsor so that they can make the decision.

Rule Structure Sub-hypothesis:

Rules can be structured to take advantage of the commonalities between satellites of different constellations.

This thesis gives support for the rule structure hypothesis by developing a generic rule structure which operates successfully on multiple satellites. A step toward verifying the main hypothesis is to present a rule structure that takes advantage of satellite commonalities. Showing support of the two sub-hypotheses will combine to give stronger support of the main hypothesis.

1.7 Approach


The approach used in this research effort began with the development of a miniature satellite expert system containing both generic and satellite-specific knowledge bases. The miniature model utilized multiple satellite-specific databases to help in its diagnosis. Once the initial GISMO prototype was developed to operate on the initial satellites, another satellite was added to the prototype to test the validity of the prototype.

1.8 Scope


This thesis focused on two communication satellites from two different constellations, the Defense Satellite Communications System Phase II (DSCS II) and the DSCS Phase III (DSCS III) pictured in Figure 1.3. These satellites were chosen based on the availability of system experts and overall research support. Designing a generic system to detect anomalies on these two satellites is one step toward showing the generic system can operate on many more satellites.

The differences between the two satellites systems are extensive. The satellites were built by different companies: TRW built the DSCS II and General Electric (GE) built the DSCS III. The two satellites use completely different methods of attitude control: the DSCS II is spin stabilized, while the DSCS III uses momentum wheels to maintain 3-axis stabilization. DSCS II is based on decades-old technology while the DSCS III hosts a modern

16-bit computer in its design. The two satellites do have a common mission and payload capability, but the payload (communication package) is not the focus of this research.



DSCS II



DSCS III

Figure 1.3 The DSCS II and DSCS III satellites

Most satellites contain the following basic subsystems: Thermal, Power, Attitude Control, Telemetry Tracking and Commanding, Mission, and Structure. To make the problem more manageable, the Telemetry, Tracking, and Commanding (TT&C) subsystem of the DSCS II and DSCS III satellites was chosen to be used as the problem domain for this research. The TT&C subsystem allows communications between the ground operators and the satellite. It accepts commands from the ground operators and performs the required actions for those commands, continuously transmitting state-of-health telemetry data down to the ground operators for analysis. The generic expert system prototype accepts TT&C

telemetry as input and detects anomalies existing within the TT&C subsystem, independent of which of the two satellites is being monitored. The generic expert system is supplemented with satellite-specific knowledge bases to detect anomalies specific to a particular satellite, though the "Proof of Concept" lies in the quality and usefulness of the generic knowledge base.

1.9 Executive Overview

This research work has shown a generic satellite monitoring expert system can diagnose anomalies on different satellites. The GISMO prototype was designed to diagnose TT&C anomalies for the DSCS II and DSCS III satellites. Once a workable prototype was in place, the Global Positioning System (GPS) satellite was analyzed to test the ability of the GISMO prototype to detect TT&C anomalies on a third satellite.

A significant number of commonalities were extracted from the DSCS II and DSCS III TT&C subsystems and incorporated into the generic knowledge base. The generic knowledge base has proven its value by reducing the number of rules required and eliminating the need for a complete satellite-specific monitoring system.

The following chapters will detail the work of this research beginning with a literature review summary in Chapter II. Chapter III contains a high-level description of the research methodology and Chapter IV follows with a detailed explanation of the GISMO design. Afterwards, Chapter V details the results of the initial two-satellite prototype and also the results of adding the third satellite to the existing prototype. Chapter VI will summarize the entire research effort and conclusions, and suggestions for future research in this problem domain.

II. Literature Review

2.1 Background

In the early to mid 1980's, a need for a more efficient satellite command and control system was recognized. Studies were performed that concluded the feasibility and benefits of incorporating Artificial Intelligence technologies into satellite command and control. In response to the need and these studies, several corporations began research and development of intelligent satellite controllers. These controllers were designed to control a specific satellite constellation.

Currently, there exist dozens of intelligent satellite controller prototypes, each designed for a particular satellite constellation with little to no capability to expand for control of other constellations. Today a new need has arisen in addition to the still existent need for a more efficient satellite command and control system. The new need is to reduce costs. In all aspects of military spending, the early 1990's is the era of reductions. The military must become more efficient and reduce costs.

As a result of this new need coupled with the existing need for efficiency, Phillips Laboratory developed the concept of the generic intelligent satellite controller. The generic controller would be able to command and control multiple satellite constellations, therefore, eliminating the need to design a separate intelligent controller for each satellite constellation. This is a new concept and little research has been performed in the area of generic intelligent satellite controllers.

It is important to understand the major work done in the field of intelligent satellite controllers over the past decade because the concept of a generic intelligent satellite controller is intertwined with the efforts of the basic intelligent controller.

This literature review examines research in the field of intelligent satellite control and some Artificial Intelligence (AI) reasoning methodologies potentially useful in this

field. The discussed research includes some systems in use today, prototype systems, and efforts no longer funded. Researchers' views on case-based, rule-based, and model-based reasoning methodologies will be examined, concluding with a discussion on the benefits of mixed reasoning methodologies.

2.2 Research and Development Programs in Intelligent Satellite Control

A significant amount of work has been accomplished in the area of intelligent satellite control. Much of the research studied and examined in this chapter began in the mid 1980's. Today, many of those systems are operational or still in development, and some have been abandoned. Most all of the research efforts are based partially, if not solely, on rule-based reasoning, while others find good results using model-based reasoning techniques. In addition, there are those who believe the answer to designing a successful satellite controller lies in capitalizing on the strengths of multiple reasoning methodologies to overcome individual weaknesses. The desired result and requirements are the major driving forces in the design of these systems.

Very little literature was found in the area of "generic" satellite controllers. Even though most of the systems discussed are built for one specific satellite family, much has been, and continues to be, learned from studying the design, architecture, program management, and methodology from inception through development of these non-generic systems.

2.2.1 StarPlan/Paragon. "STAR-PLAN is a study of the architecture and knowledge representation required to deal with diagnostic and planning tasks encountered in the operation and maintenance of satellite systems" (6:274). StarPlan is a model-based diagnostic prototype for satellite control, developed by the Sunnyvale Division of Ford Aerospace Corporation, now known as LORAL Space and Range Systems. The program

began in 1983 when Ford Aerospace was contracted by Rome Laboratories to build a prototype satellite intelligent controller to detect and resolve satellite anomalies. Rome Laboratories supported the program for nearly one year until government funding was cut. After Rome Labs terminated their funding, Ford Aerospace decided to continue the work under their Independent Research and Development (IR&D) program. They funded the program until 1990 when they determined there was not enough interest in the field to continue the program.

From the outset of the program, a major goal of the StarPlan system was to build a generic mechanism that would act on satellite-specific data allowing quick and efficient generation of new expert systems for different satellites. StarPlan was designed to evolve through three generations. The following discussion will step through the first two generations of StarPlan, including architecture design and lessons learned. The third generation was never completed because of the funding cut. This section will conclude with a discussion of the unachieved goals of the third generation.

2.2.1.1 StarPlan I. The first generation was strictly a rule-based system, initially built with the OPS5 expert system shell and later converted to the Knowledge Engineering Environment tool (KEE) from Intellicorp. The StarPlan I architecture consisted of five major components: Guardians, Monitors, Meta-Monitors, Data Bases, and a Simulator. The Guardians analyzed incoming satellite telemetry data and derived a set of hypotheses for all unexpected telemetry values. The Monitors operated over a certain class of possible anomalies, and used the hypothesis list from the Guardians to reason about the solution to an anomaly. Meta-monitors controlled the interactions of individual monitors. The Databases held facts and other data required by the components of the system and the simulator modeled satellites for testing purposes.

Many lessons were learned throughout the StarPlan I design and implementation process. With respect to system structure, researchers found it very costly and difficult to manage contextually partitioned rule sets. In addition, StarPlan ran into difficulties detecting a single anomaly that affected components in more than one object class because multiple monitors would be activated and the meta-monitor was not designed to conclude the possibility that only one fault had caused the entire problem. Each alarmed monitor would, unsuccessfully, try to solve its individual problem. This problem meant the StarPlan design could not resolve a failure whose symptoms existed in multiple object classes. Finally, the syntactic structure of rules limit the possible reasoning operations to simple pattern-matching techniques, prohibiting the development of a generic reasoning mechanism. As stated earlier in this section, a major goal of this project for Ford Aerospace was to design a system that could operate on any satellite-specific data, allowing quick and efficient generation of new expert systems for different satellites. Other lessons were learned in the areas of testing and knowledge engineering, which apply to all rule-based expert systems. The first lesson learned was the difficulty of verification and validation. Researchers realized the only true means of verifying and validating such a system would require a test of every possible rule combination, a costly solution. Another lesson learned was the difficulties caused by the knowledge acquisition bottleneck.

In addition to the above mentioned problems, StarPlan researchers determined their completed project would require more than 80,000 rules in a totally rule-based system. With that, the researchers abandoned the idea of an exclusive rule-based system and began a new design for the second generation StarPlan. (4:1)

2.2.1.2 Paragon. Paragon is a tool built by Ford Aerospace to overcome the weaknesses of the StarPlan I design and to achieve their initial goal of a generic mechanism capable of generating new expert systems for different satellites. Paragon is a

one-of-a-kind revolutionary expert system development tool. It is an elaborate and complex model-based reasoning tool that can accept any satellite model and detect, diagnose and resolve anomalies. The Paragon design addresses three major areas: knowledge representation, knowledge management and knowledge manipulation.

The Paragon knowledge representation is a unique hybrid scheme composed of frame, semantic network, classification hierarchy, blackboard, demon, and transition network technologies (22:51).

The knowledge management area includes knowledge acquisition, knowledge translation, and knowledge verification and validation. The knowledge acquisition capability of knowledge management is a major capability of Paragon. Through an elaborate user interface, "Paragon allows an expert to transfer his mental model of the domain to the computer without being taxed by normal coding and software development procedures" (11:1). Once the domain expert has entered the system knowledge into Paragon, the knowledge management component translates the knowledge into the knowledge representation format. Paragon automatically creates the necessary code to build an intelligent satellite controller. After the domain knowledge is entered into Paragon, validated and translated into the internal knowledge representation scheme by knowledge management, the knowledge manipulation portion of Paragon applies problem solving modules to the model to reason about system faults.

The knowledge manipulation area consists of four reasoning modules. First, the Data Monitoring module simply checks to see if incoming telemetry values are within prescribed limits. If a telemetry point is out-of-limits, the next module, called Situation Assessment, is notified. The Situation Assessment module generates a ranked list of components that may have caused the detected out-of-limits condition. Third, the Goal Determination module begins at the top of this list and identifies goals that would return an out-of-limits component to nominal values. Last, the Planning module searches for

events that have the potential to achieve the goal(s). If one of these plans is executed, the Planning module monitors the vehicle response to insure the problem is corrected. (11:2)

This complete redesign resulted in a model-based development tool that can accept any satellite domain knowledge, convert the knowledge into a model, and reason about any problem even if the problem has never been seen before. This was not possible with the rule-base approach. Paragon not only resolved the deficiencies found in the StarPlan I design, it established a new state of the art technology in the design of satellite controllers.

2.2.1.3 StarPlan II. The second generation of StarPlan was the first application built with Paragon. StarPlan II was designed to operate on the Electrical Power Subsystem (EPS) of the Global Positioning System (GPS). When LORAL terminated the program in 1990, parts of most of the GPS subsystems had been added to StarPlan II and the program detected, diagnosed, and resolved anomalies successfully. There was one problem carried over from StarPlan I, which was the difficulty of knowledge acquisition. The researchers found the satellite design engineers were unable to consistently produce the level of detail requested by Paragon because questions were asked after the satellite was in orbit. Ford Aerospace proposed a solution to this problem: have the satellite contractor build the Paragon model during development as part of the satellite contract (9). This would ensure the satellite model was complete and accurate, and then the intelligent ground controller could be in place and ready for operations before the satellite is launched.

2.2.1.4 StarPlan III. StarPlan III was to be the final generation of this research effort, but funding cuts came too soon. The goal of the third generation was to develop a hybrid system. The hybrid would include the Paragon model-based capabilities along with a procedural system that would perform simple telemetry monitoring tasks in

search of abnormal conditions indicative of some known anomaly. This procedural monitoring system could easily be performed by a rule-based system. The monitoring system would detect known anomalies and would notify the operator and step him or her through the pre-defined anomaly resolution passplan. If, after all applicable passplans were implemented, the anomaly was not resolved, the model-based system would take over and make suggestions as to the cause and recovery of the anomaly. (9)

2.2.2 MARVEL. MARVEL is the Multimission Automation for Real-time Verification of spacecraft Engineering Link developed by the Jet Propulsion Laboratory (JPL). MARVEL was originally built in 1989 to support the Voyager Spacecraft's encounter with Neptune. JPL continues to maintain MARVEL with upgrades scheduled every six to ten months. MARVEL operates within an X-windows/Motif environment using a combination of conventional automation techniques and an embedded knowledge-based system. Together, the two systems "provide real-time monitoring of data from multiple spacecraft subsystems, real-time analysis of anomaly conditions, and both real-time and non-real-time productivity enhancement functions." (14:81) MARVEL improves satellite operations efficiency and accuracy, and reduces the need for continuous human expert support. The task of monitoring the Voyager spacecraft, at one time requiring 4.5 people is now performed by MARVEL with the assistance of one part-time supervisor to monitor MARVEL's outputs. (5:253)

2.2.3 SHARP. SHARP is the Satellite Health Automated Reasoning Prototype developed also by JPL. SHARP was designed to capture the knowledge of the lead Voyager telecommunication expert at JPL and to assist operators of the Voyager spacecraft during the Neptune encounter. SHARP was built with a Lisp-based expert system shell on a Symbolics 3670 and includes a graphical user interface built with the

DataViews commercial tool. The SHARP design combined both conventional programming and rule-based reasoning methodologies. (3:358)

2.2.4 ASW. ASW is the Advanced Satellite Workstation developed by The Aerospace Corporation. ASW began as an Independent Research & Development (IR&D) project in 1986 to evaluate the utility of using expert systems in satellite control. The ASW was designed to control the Electrical Power Subsystem (EPS) of the Combined Release and Radiation Effects Satellite (CRRES), a complex scientific satellite.

ASW was initially developed as an exclusively rule-based system, but researchers soon realized this architecture was too limited. After some study, the researchers determined human experts use many different sources of information to diagnose and resolve satellite anomalies, including diagrams, satellite handbooks, pictures taken at the factory during design and assembly, simulations, trend analysis, and experience. The Aerospace researchers set out to build a system to incorporate all this information into some form that would be easy to navigate through and utilize. The end design was a form of multi-media system. The design included a rule-base, scanned textual information gathered from satellite handbooks, computer aided design diagrams, laser disks and video tape of the design and assembly of the satellite, and an elaborate graphical user interface. The rule-base was created with the Nexpert development tool. Because operators were not comfortable with the capabilities of intelligent systems, the satellite handbook textual data was scanned into ASW to allow the system to augment every recommendation with the corresponding section of the handbook.

In addition to ASW's normal monitoring tasks, it performed prescheduling of all required experiments for forty instruments located on the CRRES. ASW also assisted operators in building satellite passplans. Both tasks were accomplished manually prior to

ASW, but now the ASW automation of these processes reduces operator workload and the number of errors.

The ASW is not a real-time system, receiving satellite telemetry only after it has been sent to the ground control center, processed into a usable format, and stored in a file. ASW uses an Ethernet connection to the control center's computer, to link to the file of stored data. This system was first demonstrated at the Mission Control Complex VI in Sunnyvale, California and now several of the ASW capabilities are being evaluated there. (1:4-1 - 4-8)

2.2.5 MARPLE. MARPLE is a model-based diagnostic reasoning tool developed by TRW Space & Technology Group's Engineering and Test Division. MARPLE uses a modified constraint suspension methodology for use with analog device models. The constraint suspension technique models all sensor data relationships and, using a system model, monitors data consistency. The constraint suspension technique is used to determine if it is consistent to believe a suspected component is the only component malfunctioning.

The designer enters the domain knowledge into MARPLE to include: system components, all inputs and outputs of the components, and the relationships (constraints) between the components. MARPLE accepts this information and builds an internal model of the system. When system telemetry (sensor) data is fed into the model, MARPLE begins to check for inconsistencies. If an inconsistency is detected, MARPLE begins to analyze the components to find the one component that caused the inconsistency. If more than one component is at fault, MARPLE builds a superstructure or group of components that could have caused the discrepancy. (13:212-213)

MARPLE was successfully used to augment an existing rule-based system built at NASA Lewis Research Center. The autonomous power expert (APEX) is the original

expert system designed to perform fault isolation and recovery for the electrical power distribution test bed of the Space Station Freedom. NASA wanted APEX to have the capability of diagnosing unanticipated problems, so they decided to add a model-based approach to their existing knowledge base (13:206-209). NASA researchers were pleased with the results. MARPLE added the capability to diagnose unforeseen anomalies which overcame the main weakness of the exclusively rule-based system.

2.2.6 SAGE and Rule Reuse. SAGE is the SATellite Ground Expert workstation developed by The Aerospace Corporation. SAGE is a prototype system built to monitor and diagnose anomalies on the Defense Meteorological Satellite Program (DMSP). Though SAGE is not a generic intelligent controller, much of the architecture of SAGE is similar to the design planned for this thesis.

In an effort to expand the SAGE expert system to operate on other satellite and launch systems, The Aerospace Corporation realized the typical expert system rule structure used to build specific satellite controllers is difficult to generalize for use with other systems. Therefore, the Aerospace researchers proposed a new approach to designing knowledge bases to allow expert systems to operate on more than one satellite. This section will discuss the SAGE architecture and explain this new approach to rule-base design.

2.2.6.1 Architecture. The architecture consists of three main components: the knowledge-based expert system, the telemetry server, and the user interface. These components operate independently of each other, allowing SAGE to take advantage of parallel processing capabilities. This independent design also allows for the reuse of the telemetry server and user interface with other satellite knowledge bases, without major redesign. The idea of reusability of systems is the primary emphasis of this

thesis. Though the direction is focused on applying the concept to the knowledge base, it is just as important for the supporting architecture to be generic and reusable as seen in the SAGE prototype.

SAGE's knowledge bases were built with the Nexpert development tool. Nexpert allows for the use of objects and rules within its knowledge representation. Telemetry points are the facts, represented as objects, and the relationships between these objects are contained within the rules. The supervisor knowledge base, subsystem knowledge bases, and cleanup procedure knowledge base form the expert system used by SAGE. The supervisor detects the anomalies and determines which subsystem knowledge base is required to solve the anomaly. Afterwards, the cleanup procedure knowledge base notifies the operator of any additional actions required to complete anomaly resolution actions. Of course, all of these actions are based on the values of the satellite telemetry being received from the telemetry server.

The telemetry server is a Fortran program that accepts raw satellite telemetry data directly from the satellite or from stored data files and transforms the data into a usable format. To reduce the processing load, the telemetry server only passes requested telemetry mnemonics and values to the expert system or the user interface. The expert system and user interface will request certain values as required. The telemetry server only passes the value if the value has changed since the last time it was requested by a particular SAGE component, reducing the processing load of the other components in the system. The ideas of separate telemetry processing and if-changed telemetry passing will be adopted in this thesis.

Last, the user interface presents the satellite data and expert system recommendations to the user in a format that is easy to read and comprehend. The interface is built using a commercial tool called DataViews. The user interface can present the data in graphical or alphanumeric format. There are circuit diagrams, graphs, charts,

and message windows used to communicate the status of the satellite to the operator.
(17:352)

The expert system, telemetry server, and user interface work independently, passing results and requests to each other to perform DMSP operations. As stated earlier, the SAGE design closely resembles that chosen for this thesis, but due to the tools being used in this thesis research, object-oriented programming will not be included in this phase of the project. Object-oriented concepts should be implemented in a later design to strengthen the reasoning mechanism.

2.2.6.2 Rule Reuse. The Aerospace Corporation is a leader in efforts to integrate Artificial Intelligence (AI) capabilities into satellite command and control. Along with SAGE, The Aerospace Corporation has developed several other intelligent controllers designed for a specific satellite or launch system. After designing SAGE, the Aerospace researchers tried to extend the SAGE rule-base to operate on a different system. During this attempt to extend SAGE, the researchers determined the traditional rule-base design is restrictive and inflexible. As a result, researchers Dr. Scott Turner and Dr. Patricia Mangan wrote an article on Reusable Expert Systems proposing a new approach to rule-base design (18:780). In this article, the two researchers suggest the traditional rule-base design is inflexible because it intermixes general and specific knowledge. To build a reusable rule-base, the authors propose, general knowledge must be extracted and separated from specific knowledge. The generic rule-base, developed with the general knowledge, would be based on a model-based representation of the problem domain. The specific rule-base would contain specific knowledge to relate the particular problem to the generic rule-base. (18:780-781)

The model-based representation of the generic rule-base for the satellite domain would contain general abstract knowledge on satellite components to include the

components functions, structure, sensor locations, and relationship to other components. Included within the generic rule-base is knowledge of how to detect and troubleshoot anomalies. The specific rule-base would contain information to determine which components represented in the generic rule-base existed in the particular problem domain. The authors refer to this specific rule-base as the translation knowledge base and the generic rule-base as the core knowledge base. (18:781)

Turner and Mangan claim the new rule-base design allows systems to operate on more than one problem domain. This advantage overshadows the fact that building the system with this new generic design methodology takes a considerable effort up front, but the benefits outweigh this initial time-consuming cost. The authors state the separation of general and specific knowledge improves the debugging, verification and maintenance tasks while creating a smaller, simpler and more understandable expert system. (18:781-782)

This research effort is more directly related to the work of this thesis than the other discussed systems, because the ideas proposed by the researchers at The Aerospace Corporation parallel the rule-base design used for this thesis.

2.3 AI Reasoning Methodologies

Throughout the literature readings, there was a general consensus between the authors on the advantages, disadvantages and appropriate applications of different reasoning methodologies. It is important to this research not only to study existing systems, but also the views of leading researchers in the field of intelligent satellite control. The views of several leading researchers in the areas of case-based, rule-based, and model-based reasoning methodologies are summarized below.

2.3.1 Case-Based Reasoning. This methodology is based on the belief that human experts use their knowledge of previous, similar experiences to solve new problems. Case-based reasoning stores the details of previous anomalies along with the required resolution actions to help resolve future anomalies. If an anomaly occurs, the collection of previous cases is searched to find the one that most closely matches the existing problem. If an identical match is found, the associated recovery procedures are implemented to resolve the anomaly. Otherwise, the case that is the closest match to the existing problem is modified to adjust for the differences and is then implemented. Case-based reasoning is a type of shortcut method to resolving problems by anticipating new problems based on prior cases. (15:5-3)

There are a few problems associated with case-based reasoning. First, the required storage for information on every problem ever encountered could grow very rapidly. Organizing, maintaining and searching such a large amount of data could be very labor intensive. Another concern is how to choose the case that is “most closely” related to the existing one. Last, it is a difficult task to modify the recovery procedures of an old case to resolve an existing anomaly. These problems must be addressed and dealt with if case-based reasoning methodology is used in a system. Case-based reasoning is best used in domains where previous cases closely resemble future ones, such as legal, medical diagnosis and mathematical applications (15:5-3).

2.3.2 Rule-Based Reasoning. Rule-based reasoning systems use “if-then” rules to represent domain knowledge. The preconditions, or “if”, part of a rule are checked against a set of facts known as the “state-of-the-world”. If the facts found within the “state-of-the-world” satisfy the preconditions of a rule, the “then” part of the rule is executed. The “then” part, once executed, could perform some action or it could alter the

current "state-of-the-world," satisfying yet another rule. The process continues until a goal is reached or no more rules can be satisfied.

Rules are the most popular knowledge representation technique in artificial intelligence (15:5-3). Rule-based reasoning has proven its abilities in many applications and is becoming accepted as a new and successful technology. Rule-based reasoning is popular for several reasons. Human experts' knowledge fits the stimulus-response form of a rule making it easy to code the expert's expertise. Small rule-bases are easily augmented with new knowledge. Another very popular capability of rule-base systems is their ability to explain the method by which they derived their conclusions. This capability is important for a new technology not trusted by the general public. The success of rule-based systems has helped to relieve some of the skepticism and strengthen the support for the field of artificial intelligence.

Even though rule-based systems are very popular, they are not without weaknesses. Large rule-bases can incur high operation and maintenance (O&M) costs, are difficult to verify and validate, are very brittle, and have limited generic processing capability. As stated in the above paragraph, an advantage of rule-based systems is their ease of knowledge augmentation; however, this advantage only holds for small rule-based systems. When rule-based systems become large, the ease of augmentation reverts from an advantage to a disadvantage. To augment a knowledge base, the programmer must know and understand the existing rules to make meaningful and accurate updates to the rule-base. On a large system, this can be difficult for a single programmer. For example, the EXCON expert system at the Digital Equipment Corporation (DEC) requires one maintenance person for every 500 rules, and there are 3500 rules in the system (8:2) (9). Maintaining this level of programming personnel can be costly. Also, as rules are added, there can be side effects, due to rule interactions, to deal with. These side effects may not

be anticipated and will go undetected unless identified through the verification and validation process.

Rule-bases are very difficult to verify and validate. Every possible combination of rules would require testing to truly validate a rule-based system. As new rules are added to the system, the knowledge base must be revalidated. Since testing every possible rule combination is not feasible for large knowledge bases, many of these systems must undergo a long test period where they are operated simultaneously with the system they will eventually replace. This process can take a very long time which is not appealing to the customer.

A weakness associated with any rule-based system is brittleness. Rule-based systems do not function well outside of their immediate domain. They are unable to handle these situations because the knowledge contained within their rules is implicit. There is no explicit, causal, or underlying knowledge about the existing domain. The knowledge is topical and limited only to the information about the domain. The rule-base system cannot be expected to, and will not, handle unforeseen circumstances and those situations falling outside of the domain of the rule-base.

According to Dr. Marilyn Golden of LORAL, another weakness of rule-based systems is the lack of a consistent definition within rules. She states that there is nothing to dictate the way the knowledge is represented within the rule. With respect to intelligent controllers, both the knowledge about how to diagnose a problem and the knowledge about the system to be diagnosed are intermixed throughout the rules. This makes it difficult for a system to distinguish which kind of knowledge it is working with. Therefore, all knowledge must be treated the same and the kind of problem solving mechanisms acting on the data is limited to pattern-matching forward and backward chainers. This weakness limits the use of generic processing mechanisms. Golden

suggested a solution to this problem: "if rules contained semantic as well as syntactic structure, more powerful, generic problem-solving mechanisms could be employed." (8:2)

The views of Golden and her peers at LORAL Space and Range Systems raise some concerns over the feasibility of a generic intelligent controller designed with a rule-based system. Golden's claim about the limitation of generic processing are to be considered, but this is a very limited area of research. Golden and her fellow researchers at LORAL are the only ones found in this literature search to have even considered the idea of generic processing using rule-based systems. This research, therefore, did not feel that this claim of the generic processing limitation of rules was conclusive.

The rule-based reasoning methodology was chosen as the tool to be used in the development of the intelligent controller prototype of this research effort. The sponsor of this thesis was interested in determining the feasibility of using a rule-based system for satellite anomaly detection. As mentioned in Chapter I, Phillips Laboratory's concept of a multimission advanced ground intelligent controller is based on the success of a generic architecture that can operate on different satellites of different constellations. With the popularity and community acceptance of rule-based systems, Phillips Laboratory requested the use of rule-based reasoning for this research effort to learn the advantages and disadvantages of such an approach.

2.3.3 Model-Based Reasoning. This problem solving methodology uses an explicit model of the domain to operate on. The model contains information about each component of the system domain, including how each component behaves and the relationship of each component with other components of the system. Model-based systems commonly use causal reasoning, reasoning from first principles and reasoning from the principle of locality to resolve a problem. (15:5-4)

Causal reasoning is based on the knowledge of how one component affects the behavior of other components. "First principles" refers to the laws of physics and mathematics. The system uses these laws to determine the behavior of the system. The model-based reasoning mechanism considers the physical connections of the system, through the principles of locality, to resolve any problems. All of these methods, along with the explicit domain model, form a problem-solving mechanism with deep domain knowledge. The model-based approach is very powerful, yet it is not without weaknesses. Without rules to supplement and add heuristics to the model-based approach, the purely model-based system could take excessive amounts of time to resolve an anomaly. Model-based reasoning should be applied to systems where the degree of experience with system anomalies is limited, and to those in which the system is required to detect and diagnose unforeseen problems. (15:5-4)

2.3.4 Blending Different Reasoning Methodologies. Captain James M. Skinner is an advocate of the benefits of blending different reasoning methodologies. He proposes that a synergistic benefit is gained from mixing shallow and deep knowledge reasoning methods. Case-based and rule-based systems are examples of shallow knowledge, while model-base systems represent a form of deep domain knowledge. Conventional programming techniques can also offer added benefit to a diagnostic system. Numerical problems and others solved algorithmically are best solved with conventional programming. (15:5-4) Skinner built the Synergistic Reasoning System (SRS) prototype to prove the value of his blending concept (16:95,96).

The SRS intelligent controller prototype simulated fault detection and correction for the Reaction Wheel Assembly (RWA) of the Hubble Space Telescope. SRS uses a modified blackboard to control the problem-solving process performed by the case-based, rule-based, and model-based reasoners, along with some conventional programming. The

blackboard is designed to dynamically switch between reasoning systems to solve a problem. The SRS was tested on a scenario that demonstrated the benefits of the synergism created by blending different methodologies. The scenario test results supported the synergistic claim. The different reasoning methodologies contributed to resolving the problem when it was appropriate. At times when one methodology was unable to resolve a particular problem, another method was able to contribute and solve the problem. The SRS prototype was very successful in diagnosing the faults on the RWA. "The combination of the paradigms provides an ability to employ historical, experiential, procedural, causal, and structural knowledge during a problem-solving session and thus enables SRS to solve all problems solvable by any of the four reasoning methodologies individually" (16:99). (16:90-99)

2.3.5 Summary. The major AI reasoning approaches have been described, including their strengths, weaknesses, and appropriate applications. It appears that no one method can solve every problem, but maybe a combination of several methods can get closer to achieving the objective. Although this idea of blending reasoning methodologies is not a well-studied area of artificial intelligence, most of the prototypes studied employ a mixture of reasoning methods. This thesis is interested in this concept as a future addition to this research.

2.4 Conclusion

This chapter discussed several of the intelligent controller prototypes designed over the past decade. Many more exist that were not mentioned here, but all are similar in their design goals and methodologies. The primary goal is to automate the current, labor-intensive tasks of satellite operations. The methodology includes the use of one or more AI reasoning methodologies and a graphical user interface to present the hard to

understand, alphanumeric telemetry data in a format that is easy to understand. The design implemented in this thesis will follow most closely to that used in the SAGE prototype. Chapter III describes the proposed design methodology.

III. Methodology

3.1 Overview

The previous chapter summarized some current research efforts in intelligent controllers. This thesis closely follows the ideas proposed by the researchers of The Aerospace Corporation. Their ideas on the benefits of reusable rules ties directly into the concept of a generic satellite intelligent controller. This thesis investigates the feasibility of building a generic satellite intelligent controller through the use of reusable (generic) rules. The success of such a system would greatly reduce the operation and maintenance costs of satellite operations.

This chapter describes the high-level methodology used during the design of the GenerIc Satellite MONitor (GISMO) expert system prototype, while a much more detailed discussion of the design and implementation process will be covered in the following chapter. The methodology used in this research effort set out to address the following questions:

- Do enough commonalities exist between satellites of different constellations?
- Can rules be structured in such a fashion to take advantage of these commonalities?
- After a basic prototype is developed, can a third satellite be added successfully?

The first question asks if enough commonalities exist. This thesis will provide the necessary data and analyses, along with recommendations to help the financial decision makers determine if “enough” commonalities exist to support future efforts to develop a generic rule-based system. The methodology used to extract the commonalities is described in Section 3.4. A rule structure was developed that capitalized on the extracted commonalities. The methodology used to develop this knowledge representation is described in Section 3.5. Once the initial two-satellite model was developed, a third

satellite was analyzed to see if the initial prototype could also operate successfully on a third satellite. This third satellite extension was also performed to provide more evidence for the viability of the generic concept. The methodology used to analyze the compatibility of a third satellite with the existing architecture is the same methodology used to extract commonalities from the first two satellites.

With the above questions in mind, an appropriate domain was needed to design a prototype around. Therefore, the next step in the research process was to choose an appropriate domain that could provide support for the generic proof of concept.

3.2 Domain

To limit the scope of this effort and to make the research more manageable, two satellites were chosen to be used in the prototype design. The Defense Satellite Communication System Phase II (DSCS II) and DSCS Phase III (DSCS III) were selected based on availability of experts and overall support. These satellites, which are similar in mission, yet different in design, are commanded and controlled at the 3rd Satellite Operations Squadron (3SOPS) at Falcon AFB.

DSCS II was built by TRW and designed with a five year life expectancy. A total of 12 successful DSCS II launches out of 16 attempts were made between 1971 and 1989. The DSCS III satellite was built by GE with a 10 year life expectancy. The DSCS III was developed as a follow-on to the DSCS II. The first DSCS III satellite was launched in October of 1982. A total of eight DSCS III satellites have been launched, and six more DSCS III satellites remain to be launched.

Once the satellites were chosen, the domain was further scoped to one subsystem common to both satellites. The Telemetry, Tracking, and Commanding (TT&C) subsystem was chosen due to a large degree of seemingly apparent commonality. The basic functionality of the TT&C subsystem is common to most satellites.

3.2.1 The Telemetry, Tracking, and Commanding Subsystem. The TT&C subsystem is divided into the uplink and downlink sections. The uplink section allows the satellite to be commanded and controlled, while the downlink section transmits satellite position (range), and health and status information. For satellites designed to operate on the Air Force Satellite Control Network (AFSCN), there is a basic functionality of the TT&C subsystem common between these satellites. This may not be true for those satellites designed around a dedicated ground antenna support segment.

Before developing the initial prototype, a study of the functionality of the two satellite subsystems was performed. This functionality is described in Appendix A. After learning the operation of the individual subsystems, a comparison was made between the DSCS II and DSCS III TT&C subsystem operations.

3.2.1.1 Comparison of the TT&C Uplink Sections. The DSCS II and DSCS III TT&C uplink sections are similar in structure, yet the telemetry data transmitted by the satellites to represent the status of this structure are where most of the differences lie. A detailed description of the two TT&C uplink sections can be found in Sections A.1 and A.2. Both satellites have two receivers, two uplink tone detectors, two decryptors and two command processing units. Also, for certain components on each satellite, there is the ability to cross-strap in-line components. The ability to cross-strap means a primary unit of one component can be configured to connect to the secondary unit on another component and the same logic holds for the secondary unit. This cross-strap capability, which components have the capability, and what is required to access the capability will become important in later discussions. Although there seem to be several similarities between these two systems, the functionality of the components and their associated telemetry vary.

The function of the DSCS II signal conditioner is identical to the function of the AMSYNC performed within the receiver of the DSCS III vehicle. There is no cross-strapping between the receiver and the signal conditioning function of the DSCS III satellite like there is for the DSCS II satellite. If any of the functions performed within the receiver of the DSCS III fail, the entire receiver is lost. Another important difference is the cross-strapping capability between the receivers and INYs of the DSCS III satellite, which is accessible by ground command only. There is no cross-strapping between the signal conditioners and the INYs of the DSCS II. Therefore, if a signal conditioner fails, the whole command path on the DSCS II is lost, whereas on the DSCS III it is not.

There are also some differences in the telemetry of the two TT&C uplink sections. Beginning with the receivers, the DSCS II receiver telemetry value is "LOCK" if either receiver detects its uplink signal, while the DSCS III receiver reads "A" or "B" depending on which receiver detected the signal. The DSCS II has receiver converter analog telemetry data that represents the voltage of the receiver converter and the DSCS III does not. The DSCS III has analog voltage telemetry to represent the status of its INY, where the DSCS II has a discrete ON/OFF/ONE/TWO representation for the status of its INY. In addition to these analog verses discrete differences in the representation of the vehicle status, there is a slight difference in how the two systems distinguish between their primary and secondary components. The DSCS III satellite uses an A to represent a primary component, while the DSCS II satellite uses ONE for the same purpose. And likewise, DSCS III uses a B to represent a secondary component and DSCS II uses a TWO. These differences in the type and format of telemetry data representing each component added to the challenge of designing a generic expert system.

3.2.1.2 Comparison of the TT&C Downlink Sections. A detailed discussion of the DSCS II and DSCS III TT&C downlink sections can be found in Sections A.3 and A.4, respectively. Much like the uplink sections, the downlink TT&C sections are also similar in structure, but different in the telemetry used to represent their status. Both DSCS II and DSCS III have redundant multiplexers, encoders, encrypters, subcarrier and carrier generators, and transmitters. Just as seen in the comparison of the uplink sections, the telemetry representation of these components vary. For example, DSCS II uses discrete ON/OFF/"selected" status representation for its encrypters while DSCS III uses analog encrypter converter voltage telemetry to represent the state of its encrypters.

A fundamental difference that occurs in both the uplink and downlink sections of the two satellites is the process required to enable and disable cross-strapping between components. Cross-strapping between components on DSCS II occurs automatically. For those components that are cross-strapped on DSCS II, there exist both a cross-strapped link and a direct link. This means the uplink/downlink signal is routed to the primary and alternate component that follows it, but only the component that is selected and powered accepts and processes the signal. This is different from the DSCS III satellite in which most cross-strapping is enabled by ground command only. The uplink/downlink signal on the DSCS III satellite is routed to only the selected unit instead of both. The disadvantage to the command-enabled cross-strap is that it adds complexity which can reduce reliability. Due to these reliability issues, operating procedures only recommend commanding the reconfiguration of some physical component after all other possibilities have been exhausted. This is true for both satellites.

Once the domain was selected, the appropriate development tools were chosen. The tools were chosen to fulfill the needs of this research as well as the needs of the sponsor.

3.3 Development Environment

Because this thesis explores the aspects of creating a generic expert system, most of the research effort was directed toward gathering data to support the practicality of a generic expert system for satellite anomaly resolution. The prototype was developed to generate ideas and to refine the design of a generic rule-base, not to design a commercial-quality satellite anomaly detection system.

The prototype was built with the EXSYS rule-based expert system development tool, written by EXSYS, Inc. of Albuquerque, New Mexico. The software was purchased by Phillips Laboratory, the sponsor of this research. In addition to EXSYS, the Paradox database software, developed by BORLAND, was also purchased by Phillips Laboratory for use in this thesis. The Paradox database includes a script language called ObjectPal, that was used extensively as an interface tool between database and expert system operations.

Rule-based reasoning was chosen for this thesis because the long-range goal of Phillips Laboratory will be to create more satellite autonomy by embedding the rule-base on the satellite. After this system has proven itself on the ground in day-to-day satellite operations and has passed all verification and validation testing, it will be moved to the satellite. Rule-base systems have been successfully used over the past decade in many different domains and this success has generated increased confidence in this AI technology. Model and case-based systems are not as well-understood and have not gained the recognition that rule-based systems have.

Rule-based reasoning is capable of resolving expected problems. It is not able to resolve those problems that are unanticipated and have never been encountered before. Again, the rule-based approach was chosen because the end user of this generic satellite controller does not require the system to handle unexpected problems. If this became a

requirement, another reasoning methodology, such as a model-based system, would be necessary to resolve unforeseen problems.

A database was required to manage the large amount of data required of a generic satellite controller. The Paradox database was chosen because of its ability to interface with the EXSYS software and because of its easy-to-use window environment.

These development tools were quite sufficient for the work done in this thesis, but as this research continues, the initial development tools will need to be extended or be replaced with more powerful tools.

3.4 Extracting Commonalities

To compare the two chosen satellites, I studied the basic functionality, normal operational procedures, recognized subsystem failures and telemetry verifications of both TT&C subsystems. I performed satellite operations for the DSCS II satellite prior to my AFIT assignment, which helped in my study of the DSCS II TT&C subsystem. Captain Bob Costa, currently assigned to AFIT, had performed satellite operations on the DSCS III satellite at his previous assignment to the 3SOPS and so provided the domain expertise for my DSCS III analysis.

All necessary telemetry, operational pass plans, and telemetry limit information was provided by the personnel at the 3SOPS. Captain Costa provided all of the Orbital Operation Handbooks (OOH) and operations manuals used by the satellite operators at Falcon AFB to operate and control the DSCS III satellite. Captain Costa was a certified trainer and evaluator of the DSCS III satellite while at the 3SOPS, and was extremely knowledgeable and helpful in providing information on the DSCS III TT&C subsystem throughout this thesis effort.

After the knowledge acquisition process was accomplished, the telemetry from each component of both TT&C subsystems was compared to extract common telemetry

points. A common telemetry point was one that portrayed the same information about a component found on both vehicles. After a telemetry point was found to be common between the two satellites, the telemetry variable name used by the specific satellite was renamed to a new generic name that represented a common telemetry point. An example can be found in the tone detector telemetry of the two satellites. The telemetry point for the DSCS II tone detector is called SIGCON which stands for signal conditioner. The DSCS III tone telemetry is AMSYNC, meaning the detector has detected an AMSYNC (modulated) uplink signal. Satisfied that this two values represented the same information, I gave the two points the common variable name, TONES. This is because the function of the two components is to detect tones on the uplinked signal.

After the subsystem commonalities were extracted, I develop a knowledge representation to capitalize on these commonalities.

3.5 Rule-Base Design Considerations

Before building the rule-base, I studied the design and functionality of the two TT&C subsystems and the current anomaly resolution plans used at the 3SOPS to resolve any known TT&C anomaly on the two satellites. The anomaly recovery passplans not only gave the recommended resolution to a particular anomaly, but it also gave the symptoms that would result from the particular anomaly. These anomaly passplans provided valuable knowledge and were very useful in developing the rule-base.

With an understanding of the basic system functionality and all of the known and well-understood anomalies that could occur on these systems, I began to build the rule-base. The rule-base was designed around the anomalies. The diagnosis of any anomaly that required sensor data common to both satellites to detect, was coded into a generic rule. If both satellites had an anomaly that required some data common to both satellites and some data specific to the satellite to diagnose, a generic rule was used to test the

common telemetry values and, if found to be true, would set some intermediate flag. This intermediate flag would cause a satellite-specific rule to be tested. The satellite-specific rule would check some satellite-specific telemetry and, if found true, would conclude that an anomaly had occurred and perform the necessary recovery procedures. An example of this is shown below.

GENERIC RULE:

IF common_sensor_data_A = some value indicative of a failure
 and common_sensor_data_B = some value indicative of a failure
 and other preconditions are met
THEN A possible component_A failure has occurred

SATELLITE SPECIFIC RULE:

IF This is Constellation [Constellation name]
 and A possible component_A failure has occurred
 and specific_sensor_data_C = some value indicative of a failure
THEN component_A has failed
 and make call to recovery procedures
 and post any required changes to the satellites expected status

Embedded variables were also used throughout the rule-base to aid in the development of more generic rules.

The prototype design was based on several simplifying assumptions. These assumptions are listed below and discussed further in Section 5.2.2.

- No multiple point failures exist.
- Telemetry flows in an uninterrupted stream.
- Anomaly recovery procedures are accomplished without error and without deviation from the original plan.
- The satellite response to the performed recovery procedures is immediate.

- All failures are hard failures. No component has the ability to fall below the lower red limit and later return to a nominal status.
- All satellites are designed to operate on the Air Force Satellite Control Network (AFSCN). (further discussed in Section 3.2.1)

Once a design structure was in place that could operate sufficiently on the TT&C subsystems of the selected satellites, the TT&C subsystem of a third satellite was analyzed and tested to see if it could fit into the existing prototype structure. This was done to test the viability of the prototype and the feasibility of a generic satellite controller. The results of this test are provided in Section 5.3.

3.6 Summary

This chapter discussed the selected domain and development tools along with the methodology used to answer three questions fundamental to the success of this thesis. First, the methodology used to address the question of whether enough commonalities exist between satellites was discussed. Though this research cannot make such a financial decision, it can make recommendations and show support for or against the development of a generic satellite controller. Next, the methodology used to develop a generic rule-base capable of taking advantage of the commonalities found in the two satellite subsystems was described. Last, a third satellite was analyzed to test the compatibility of the prototype using the same methodology as that used for the first two satellite subsystems. The following chapter discusses, in much more detail, the architecture design of the prototype.

IV. Design and Implementation

4.1 Overview

The previous chapter discussed the methodology used in the design of the GISMO prototype, the scope of the problem, and the development tools used in the design. This chapter will give a detailed description of the GISMO architecture, including the supporting infrastructure as well as the rule structure. But first, I will explain why such an extensive I/O interface was built for the prototype and how GISMO will fit into the big picture of MAGIC.

4.2 MAGIC, TAS, and GISMO

Prior to the start of this thesis, MAGIC was only a concept. No software had been developed. But in the middle of my research, my sponsor was tasked by the 3SOPS to build a Telemetry Analysis System (TAS) to help ease their conversion from operators trained on specific satellites to generic trained operators. The TAS system was requested with a delivery date of October 1994. Therefore, while I was building my generic prototype which required an I/O interface to support the generic rule-base, my sponsor was building an elaborate I/O system for the 3SOPS. The TAS system is the beginning of what will one day be MAGIC. The TAS system is strictly a telemetry limit checker that offers additional easy-to-understand information to the operator such as plots, graphs, or history data. No generic concepts or anomaly detection capabilities are incorporated into the TAS system at this time. This thesis will add some insight into the issues that will need to be addressed to make a generic system successful, easing the transition process from TAS into MAGIC.

Because the development of GISMO began prior to TAS, a temporary I/O interface was developed to handle the data-intensive needs of the prototype expert system. Due to the limitations of the tools used in the design of GISMO, several work-arounds were necessary to

insure proper handling of the data. These work-arounds are temporary and will be replaced with the TAS system once it is operational.

Figure 4.1 is a pictorial representation of the GISMO prototype architecture. This architecture is a subset of the MAGIC architecture described in Section 1.3. The functions performed by TAS are simulated in GISMO through the use of Paradox scripts. Also, the interactive Graphical User Interface (GUI) anomaly recovery system planned for MAGIC is simulated in GISMO with text files. These necessary work-arounds used in the GISMO design are described below in detail, along with the capabilities of some new software tools planned for use in future efforts.

4.3 Work-Arounds

The GISMO expert system requires large amounts of external data to perform its diagnostic functions. The expert system requires the values of the satellite telemetry, the values of the expected satellite state, and satellite ephemeris data. All of this data is stored in satellite-specific databases. Each satellite has its own telemetry database which simulates the telemetry being transmitted from the satellite. In addition to the telemetry database, each satellite has its own limits/status database to hold the expected state of the satellite and an ephemeris database to hold satellite position information. This is a very intensive I/O process. To obtain a database value, EXSYS opens the database, reads one value, and closes the database. This becomes a time-consuming problem when the expert system requires over a hundred values and it must open and close the database for each value read. To help reduce the work load of the expert system, Paradox scripts were written to read the data from the satellite-specific databases and write them into a text file in a format usable by EXSYS. Therefore, instead of reading one value at a time from the databases directly, the expert system quickly reads all of the variable values at once from the text files created by the Paradox




Figure 4.1 The GISMO Prototype Architecture

scripts. The TAS system will include a direct Dynamic Data Exchange (DDE) link from the EXSYS rule-base to the databases which will eliminate the need for the Paradox scripts.

Another necessary work-around used in the GISMO design was the use of text files to simulate recovery procedures. EXSYS and expert systems in general are not designed to operate in a dynamic, changing, environment and Microsoft Windows does not handle multitasking environments very gracefully. Because of these limitations, GISMO was not designed to interact with the operator and respond to the operator's actions. Instead, the prototype embodies the assumption that whatever is recommended by the expert system is carried out by the operator without error and the satellite responds immediately and as expected.

Because of the above assumptions, GISMO simulates the anomaly recovery through the use of text files. After recovery from an anomaly, the expected satellite state may change and this change will need to be reflected in the limits/status database. This update to the expected state is simulated with the use of external text files. The text files represent the call to a recovery procedure, and they contain the necessary expected variable updates in response to a successful recovery implementation. After the detection of an anomaly, GISMO simply appends a text file, containing the required expected state changes, to a running update file. At the end of a rule run, the updates are made to the limits/status database by a Paradox script.

To overcome these current limitations, the sponsor has purchased new software for future research efforts. This software is described in Section 5.2.4.1. Until these new capabilities are integrated, the work-arounds are necessary to insure proper operation of the prototype. To better understand the operation of GISMO, the scripts used to interface the expert system with the satellite-specific databases will be explained, but first the format of the databases will be described.

4.4 Supporting Databases

There are several databases required by the GISMO expert system to properly detect satellite anomalies. The telemetry database holds the processed telemetry transmitted by the satellite. The limits/status database contains the limit ranges for the analog telemetry points along with the status of discrete telemetry. Last, the ephemeris database holds satellite position information with respect to each visible ground antenna. The structure and purpose of these databases are described in greater detail below.

4.4.1 The Telemetry Database. Each satellite has its own telemetry database. Each record within the telemetry database represents consecutive time slices of a satellite contact. A small excerpt of the DSCS II telemetry database is shown below in Figure 4.2.

Rec#	Time	Invs	Cmd1ck	Cmd2ck	Rc1p15	Rc2p15	Mxcalv	Encdr	Encalv	Encrp	Eiamat
1	45,613.91	OFF	PL	VERF	0.05	14.74	3.31	TWO	3.31	ON2	BOVR
2	45,615.51	OFF	VERF	VERF	0.05	14.73	3.31	TWO	3.31	ON2	NORM
3	45,617.16	OFF	VERF	VERF	0.05	14.73	3.31	TWO	3.31	ON2	NORM
4	45,618.86	ONE	VERF	VERF	0.05	14.73	3.31	TWO	3.31	ON2	NORM
5	45,620.56	ONE	VERF	VERF	0.05	14.73	3.31	TWO	3.31	ON2	NORM

Figure 4.2 DSCS II Telemetry Database Excerpt

The column heading names are the satellite telemetry mnemonics, or renamed common variable mnemonics, and the values under the column heading are the values telemetered by the satellite for that sensor point at the corresponding time found in the time column.

4.4.2 The Limits/Status Database. Another database used in the GISMO design is the limits/status database. Each satellite has its own limits/status database. This database contains the lower red (LR), lower yellow (LY), upper yellow (UY), and upper red (UR) ranges for the floating point telemetry values. It also contains expected and redundancy status

of the discrete telemetry points. A small excerpt of the DSCS II limits/status database is shown in Figure 4.3.

<u>Rec#</u>	<u>Tlmpoint</u>	<u>LR</u>	<u>LY</u>	<u>UY</u>	<u>UR</u>	<u>Status</u>	<u>Redundancy</u>
1	CMD1CK					VERF	
2	CMD2CK					VERF	
3	EIAMAT					NORM	
4	ENCALV	3.31	3.32	3.34	3.35	WORKING	
5	ENCDR					ONE	YES
6	ENCRP					ON2	NO
7	INYS					BOTH	YES
8	MXCALV	3.31	3.32	3.34	3.35	WORKING	
9	RC1P15	14.70	15.00	15.15	15.30	WORKING	
10	RC2P15	14.70	15.00	15.15	15.30	FAILED	
11	DOWNLINK					15	
12	PRIMARY					ONE	

Figure 4.3 DSCS II Limits/Status Database Excerpt

Record #1 states the expected value of CMD1CK is VERF and no redundancy is shown because the redundant unit is CMD2CK in record #2. The lower red limit of the receiver converter #1 (RC1P15), found in record #9, is 14.70, the lower yellow limit is 15.00, the upper yellow limit is 15.15, the upper red limit is 15.30 and the status of the converter is working. For floating point variables, the status column is used to hold the failed/working status of the component since there is no expected value as there is for discrete variables. There is also no redundancy associated with a floating point variable because both the primary and the redundant units have their own telemetry data. As a last example, the encoder (ENCDR), shown in record #5, has an expected value of ONE and redundancy is available. It should be noted that a database such as this, containing the availability of redundant units, would be classified in a real-world scenario.

4.4.3 The Ephemeris Database. The last database used in the GISMO prototype is the ephemeris database. The ephemeris database holds the relative position of each ground antenna with the satellite of interest. Again, each satellite has its own ephemeris database. A sample ephemeris database for satellite 9446 is shown in Figure 4.4.

Rec	RGF	Azimuth	Elevation	Range
1	SCFIOSA	22.00	77.00	40,456.00
2	SCFGTSA	160.00	50.00	40,279.00
3	SCFGTSB	160.00	50.00	40,279.00
4	SCFHULA	182.00	33.00	40,489.00

Figure 4.4 9446 Ephemeris Database

The ephemeris database in Figure 4.4 contains the position of satellite 9446 with respect to the satellite ground antennas within visible range of 9446. Rec is just the database record number. RGF stands for Remote Ground Facility. The values in the RGF column are call names for the different satellite ground antennas. For example, SCFIOSA is the call name for the Satellite Control Facility in the Indian Ocean; SCFGTSA and SCFGTSB are the Satellite Control Facilities in Guam. There are two ground antennas at GUAM, one is given the "A" designator and the other a "B" designator. And last, SCFHULA is located in Hawaii.

In the current Air Force satellite control environment, position information is collected from a satellite every time a contact with the satellite is made. Orbital analysts process the position information using orbital mechanics techniques, and make updates to the ephemeris database to reflect any movement of the satellite with respect to the ground antennas. The MAGIC system will perform these functions automatically. For GISMO, the ephemeris data is assumed to be static for each satellite and no updates are made.

The ephemeris database is used by the expert system during situations of a loss or a lack of telemetry, during which the expert system uses the ephemeris database to ensure the ground antenna is pointing correctly.

Much of the data in the telemetry and limits/status databases is required by the expert system at the beginning of the rule run. This data acquisition process can be very time consuming; therefore, scripts were built to decrease this time. Each script used in GISMO is described in the following section.

4.5 Paradox Scripts

Paradox scripts are used to translate the values found in the telemetry and limit/status databases into a format that EXSYS can read quickly. There is also a script used to make updates to the limit/status database in response to recovery procedures.

4.5.1 The Expected State Script. This Paradox script is called by GISMO at the beginning of the run. The script opens the appropriate limit/status database, reads the necessary satellite expected values from the database and writes these values into a text file. For discrete/bi-level sensor points, the script writes the expected value and the available redundancy to the text file. For the floating point values, the script writes the failed or working status of the represented component to the text file. The variable format used in GISMO to hold these expected and redundancy status values is described in Section B.1. Once the file has been created, the GISMO expert system will read the data and initialize its variables.

4.5.2 The Satellite State Script. This script opens the appropriate telemetry and limit/status databases and creates another text file. The script reacts differently depending on whether it is reading the very first record of telemetry data or any subsequent record.

When reading the first record of the telemetry database, the script writes the variable/value pair to the text file in a format readable by EXSYS. The variable format used to represent actual satellite state data is described in Section B.1. Once the telemetry data is written, the script compares each telemetry value to its corresponding expected value found in the limit/status database. If a floating point variable is in the nominal range or if a discrete/bi-level telemetry point equals the expected value, the point is ignored (because GISMO defaults to a nominal state). If the floating point telemetry value is not within its nominal range (OOL) or a discrete telemetry point does not equal the expected value, the script sets a qualifier value to correspond to the OOL condition of the telemetry point. The qualifier format used to describe the different OOL conditions within GISMO is explained in Section B.1.

While reading the remaining records of the telemetry database, the satellite state script compares the present value of a data point to its previous value and ignores those points that have not changed in value. Those points that have changed in value are compared to the limit/status database and if they fall in the expected range or equal the expected value, their corresponding qualifier is set to nominal. If the changed telemetry point value falls outside the expected range or does not equal its expected value, the appropriate qualifier value is set to represent the new value.

This script is run by GISMO prior to every pass through the rules. Once the script creates its text file, GISMO will read the values and set its internal qualifiers and variables accordingly.

4.5.3 The Update Script. The update script is called by GISMO at the end of each run through the rule-base. The purpose of this script is to update the appropriate limits/status database with any requested changes in response to any performed recovery procedures.

4.6 Generic Rule-Base Design

Once the databases and scripts were in place supporting the I/O requirements of the expert system, efforts were focused on the rule structure and how to make it generic. Most of the existing satellite intelligent controllers, discussed in Chapter II, are designed to operate on one specific satellite constellation. The rules found in these specific satellite intelligent controllers contain explicit telemetry names and limits. For example:

```
IF      RC1P15 > 16.4
or      RC1P15 < 14.5
THEN    The Receiver Converter has failed.
```

To avoid hard-wiring the limits and telemetry point names as shown in the rule example above, I used variables, whose values would be found in a satellite-specific database, and qualifiers. This table look-up method allowed for a more generic rule structure such as the one shown below.

```
IF      Telemetry is present
and     [AACTV_PASV] = ACTIVE
and     The actual value of TONES is abnormal
and     The actual value of RCVRLK is nominal
and     [SPE] = E
THEN    [TONE_FAILED_READS] IS GIVEN THE VALUE [ATONES]
```

This is a completely generic rule. The variables [AACTV_PASV] and [SPE] and the qualifiers for TONES and RCVRLK are common to both satellites, including the status of telemetry. The values of the variables and qualifiers are derived from the satellite-specific databases.

The database values of TONES and RCVRLK are used to set the value of the corresponding qualifiers. TONES and RCVRLK are generic representations of independent components found on each satellite having the same functionality. As stated in Chapter III, when these common components are identified, the mnemonic used to represent the component on a particular satellite is renamed and given a generic name to be used by the expert system. This renaming process is a function of the telemetry preprocessing system.

To be more specific, the generic name TONES was created to replace the satellite-specific terms of AMSYNC for DSCS III and SIGCON for DSCS II. These components on the different vehicles perform the same function. They both look for tones on the uplink signal. RCVRLK is also a generic telemetry point name created to replace two specific points with a common function. SIGPR is the specific name used on DSCS II to hold the value representing the lock status of the receivers, and RCVRLK performs the same function for DSCS III. RCVRLK is the common name used by the expert system to detect anomalies. These common names contributed directly to the generation of a more generic rule structure.

In the consequence of the above rule, the variable [TONES_FAILED_READS] is given a value as a result of the preconditions being satisfied. This variable is different from the others. The value of this variable cannot be found in any database, its value is derived as a result of the firing of certain rules. This variable holds intermediate knowledge within the rule-base. This intermediate knowledge is used to trigger other rules to fire, like the one shown below.

```
IF      [TONE_FAILED_READS] = "TWO"  
  and   Other Satellite-specific Preconditions are met  
THEN   Call recovery subroutine [[Sat]]tone.rcv
```

If the intermediate knowledge variable was given the value of TWO in a previous rule; and "Other satellite-specific preconditions are met", this rule will fire. This intermediate knowledge technique is used to diagnose anomalies which may have several common preconditions that must be satisfied and also some specific ones. To reduce duplicate effort, a generic rule tests the common preconditions and if all are true, sets the intermediate knowledge variable to cue the satellite-specific rules to be tested. Any further specific preconditions are checked and if satisfied, the recovery procedure is called using an embedded variable to make the call generic.

Early in the GISMO design, this intermediate knowledge was represented using a qualifier instead of a variable. This approach was abandoned because the quality of the

intermediate knowledge was limited to a failed/working status. Qualifiers cannot be set to the value of a variable; therefore, the satellite-specific rule would need to perform an additional check to determine the status of the actual telemetry. An example of this is shown below.

Generic Rule:

```

IF      Telemetry is present
  and   [AACTV_PASV] = ACTIVE
  and   The actual value of TONES is abnormal
  and   The actual value of RCVRLK is nominal
  and   [SPE] = E
THEN A possible TONES failure has occurred.      ;; qualifier is set
  
```

Satellite-specific Rule:

```

IF      A possible TONES failure has occurred.
  and   [ATONES] is given the value TWO      ;;additional check
  and   Other Satellite-specific Preconditions are met
THEN    Call recovery subroutine [[Sat]]tone.rcv
  
```

Figure 4.5 is a pictorial representation of the current rule structure. It should be noted that there are a few anomalies which can be detected without any satellite-specific checks. Through the use of embedded variables, the generic preconditions are checked, the anomaly is detected, and the simulated recovery procedure call is performed. An example of a true generic anomaly detection is shown below.




Figure 4.5 GISMO Generic Rule Structure

RULE #1

IF: Telemetry is not present.
and The RGF is not autotracking (no carrier present).
and Is the RGF configured for downlink channel [[DOWNLINK]]? YES
and The azimuth and elevation angles found in the ephemeris match the pointing angles of the RGF antenna.
and All ground equipment has been thoroughly checked. True
THEN: [FAILED_TMX] is given the value [TMX]

RULE #2

IF: [FAILED_TMX] = [TMX]
and [TMX_RD] = "NO"
THEN: A [[CONSTELLATION]] transmitter has failed with no redundancy.

RULE #3

IF: [FAILED_TMX] = [PRIMARY]
and [TMX_RD] = "YES"
THEN: The [[CONSTELLATION]] transmitter [[PRIMARY]] has failed.
and Append expected state changes text file to the running update file

RULE #4

IF: [FAILED_TMX] = [SECONDARY]
and [TMX_RD] = "YES"
THEN: The [[CONSTELLATION]] transmitter [[SECONDARY]] has failed.
and Append expected state changes text file to the running update file

Rule #1 tests some preconditions common to both satellites and, if found true, sets an intermediate knowledge variable to represent the possible failed transmitter (TMX). Rule #2 checks to see if the first rule set the intermediate knowledge variable ([FAILED_TMX]). If the intermediate variable is set in the first rule, the second rule checks the availability of redundant transmitters. If no redundancy exists, then a generic choice is set using embedded variables. The value of CONSTELLATION (DSCS II or DSCS III) will be placed in the choice and displayed to the user. The third and fourth rules check whether the failed TMX is the primary or secondary TMX. Again, extensive use of embedded variables allowed the generic rules above to be used in place of satellite-specific ones. A specific example of the use

of embedded variables is the use of the variables Primary and Secondary. The Primary and Secondary variables are not telemetered by the satellite, they are definition variables defined in the satellite-specific databases. These variables aid in the design of generic rules, as shown in the above example.

4.6.1 Special Rule Structure Requirements. Two additional preconditions were necessary in most rules for proper operation of the expert system. The first additional precondition is required in every satellite-specific rule. This precondition tests which satellite constellation is transmitting the current telemetry. The precondition tests if "The satellite constellation is DSCS II" or if "The satellite constellation is DSCS III". This is necessary to ensure EXSYS does not attempt to find values of variables that are unnecessary for a particular support. For example:

```
IF    [AKG46BY] = "NO"  
    and Other pre-conditions are met  
THEN Recover
```

Since all rules are tested during each run of the expert system, there is nothing in the above rule to prevent EXSYS from attempting to find the value of the variable AKG46BY. The problem is variable AKG46BY is specific to DSCS III. If the expert system is running a support on DSCS II, EXSYS will not be able to find the value of this variable in the satellite state and will be forced to ask the operator for its value. This will not make sense to the operator because this telemetry point does not exist on DSCS II. Therefore, each satellite-specific rule is prefaced with the constellation precondition and if that precondition fails, EXSYS will not test any other preconditions of the rule. For example:

```
IF    The satellite constellation is DSCS III  
    and [AKG46BY] = "NO"  
    and Other pre-conditions are met  
THEN Recover
```

The additional constellation precondition would not be required if EXSYS allowed the run of selected non-contiguous rules. If this was possible, a simple rule for each satellite, similar to the one shown below, would be sufficient to ensure EXSYS only ran the rules pertinent to the satellite being contacted.

```
IF      The satellite constellation is DSCS III
THEN RUN rules 1-27
      and RUN rules 48-67
```

One group of rules would represent the generic rules and the other would be the satellite-specific rules. This is a capability EXSYS does not have, so the constellation precondition is required at the top of every satellite-specific rule.

There is another precondition necessary to handle a no-telemetry situation. Any rule requiring the value of a satellite state variable must be prefaced with the precondition, "Telemetry is Present". Again, this insures EXSYS does not ask the user for the value of the variables during a no-telemetry anomaly. The above rule would now look like the following:

```
IF      Telemetry is present.
      and The satellite constellation is DSCS III
      and [AKG46BY] = "NO"
      and Other pre-conditions are met
THEN Recover
```

To test if telemetry is present, the expert system checks the IRON satellite state value. The IRON value is the first item in the raw satellite telemetry stream. If this value is not valid, it is a good assumption the remaining telemetry is also invalid.

With the overall design described above, GISMO can identify certain anomalies and simulate recovery procedures. Following is a description of the steps taken by the GISMO expert system during a run.

4.7 GISMO Operations

When initiated, the GISMO expert system reads the values of the actual satellite state and the expected satellite state from the appropriate text files. After the variable values are loaded, the rules are run using forward chaining. If there is an anomaly in the satellite state, the rule(s) associated with the anomaly will fire and make a simulated call to a recovery procedure which appends the necessary expected state updates to an update file. After the simulated recovery, the rules continue to operate on the initial satellite state and detect any further anomalies that may exist. At the end of the rule run, all of the required changes to the expected state will be updated in the limit/status database. After the expected values are updated, GISMO will read the next set of satellite state variables and the new expected satellite state, and repeat the above process until the end of the simulated satellite contact (the end of the telemetry database is reached). Each set of satellite state variables represents a time slice from a real satellite contact.

4.7.1 GISMO's Limitations. GISMO has several limitations which should be eliminated with the addition of increased capabilities in future efforts. The current temporary limitations are listed below.

- GISMO cannot handle any deviation from the planned recovery procedure.
- All failures are considered hard failures. This means, if a component falls below its lower red limit it is considered failed with no chance of recovery.
- Multiple failures, in which the symptoms of one anomaly are hidden by the symptoms of another anomaly, are not detectable by the GISMO expert system.

- There is no memory capability in GISMO. Once a recovery procedure is performed the system must see an immediate change in the satellite state or it will diagnose the same anomaly again.
- Due to the memory limitation, any anomaly resolution that could require several component swaps before the failed unit is found is not possible under the current design.
- The system cannot handle a telemetry data hit where variable values become scrambled (noisy data).

4.8 Summary

This chapter presented a detailed description of the GISMO architecture. The chapter began with a discussion of the TAS system and how it and GISMO fit into the MAGIC master plan. The GISMO architecture was presented along with a detailed description of the databases and scripts used to support the needs of a generic and satellite-specific rule-base. Each tactic used to create the generic structure was also explained, and last, the operation and limitations of the prototype were discussed.

Through the process of designing the GISMO prototype and analyzing the possibility of adding a third satellite, many concerns and insights arose. These concerns and insights along with some personal opinions will be discussed in the following chapter.

V. Results and Issues

5.1 Introduction

This chapter is divided into two major sections. The first section discusses the analysis of the two-satellite prototype, to include the extracted commonalities, detected anomalies, and prototype limitations. The second section details the results of a significant test of the generic concept in which a third satellite was added to the original two-satellite prototype.

5.2 GISMO, the Initial Two-Satellite Prototype

This section begins with a discussion of the commonalities extracted from the DSCS II and DSCS III satellites as a result of a direct one-to-one TT&C telemetry comparison. Following the discussion of the satellite commonalities, the anomalies successfully detected on both satellites using the GISMO generic rule-base will be discussed.

5.2.1 Telemetry Commonalities. I performed a detailed comparison of the two original satellites, to include a one-to-one TT&C telemetry comparison. The detailed comparison can be found in Section C.1 and a synopsis is shown in Table 5.1.

	Total TT&C Tlm Points (X)	# Common with other satellite (Y)	Percent in common (Y/X)
DSCS II	45	18	40.00%
DSCS III	66	18	27.27%

Table 5.1 DSCS II and DSCS III TT&C Commonalities

DSCS II has 45 sensor points associated with its TT&C subsystem and DSCS III has 66. After the telemetry comparison, 18 sensor points were found to be common between the two satellites. More specifically, these 18 telemetry points relayed the same functional information about their corresponding component. These 18 points were given a common name for use in the expert system. The GISMO expert system uses several of these common telemetry points, along with some satellite-specific telemetry points, to diagnose seven major TT&C anomalies.

Table 5.2 gives a list of the common telemetry points, along with some definition variables defined within the satellite-specific databases and some intermediate knowledge variables. These are the variables used by the GISMO expert system to diagnose, in whole or in part, seven major TT&C anomalies. The detectable anomalies are: (1) failed receiver, (2) failed tone detector, (3) failed INY, (4) failed encoder, (5) failed encrypter, (6) failed subcarrier generator, and (7) failed transmitter.

Another analysis was performed to determine the percentage of detectable TT&C anomalies on each satellite. The simplifying assumptions used in the GISMO design limited its diagnostic capability. More specifically, since no commanding capability was incorporated into the GISMO prototype, no commanding anomalies could be implemented. This does not affect the results of the generic rule-base because the majority of commanding anomalies are specific to the satellite and would not be detectable using a generic rule. Sections C.3.1 and C.3.2 show the major TT&C anomalies of DSCS II and DSCS III, respectively, along with the detectable status of those anomalies using the generic rules of GISMO. All of the DSCS II major TT&C anomalies are detectable using the GISMO expert system, though this detection is only one level deep. This means the most probable cause of the anomaly is assumed to be the only cause and no other

Generic Variable	Brief Description
IRON	satellite operations number
RCVRLK	receiver lock status
TONES	tone detector status
INYS	decrypter status
ENCDR	encoder status
TMX	transmitter status
SPE	S-pulse status
ACTV/PASV	active/passive status
AZ	azimuth
EL	elevation
RANGE	range
Definition Variables Defined in the Limits/Status Database:	
CONSTELLATION	the satellite constellation name
SAT	an abbreviated constellation name used in file calls
DOWNLINK	the number of the downlink channel being used
UPLINK	the number of the uplink channel being used
TONE_OFF	the value the tone detector reads when it is totally failed
INY_OFF	the value the INY reads when it is totally failed
RCVR_OFF	the value the receiver reads when it is totally failed
PRIMARY	the symbol used to represent primary components
SECONDARY	the symbol used to represent secondary components
Intermediate Knowledge Variables:	
FAILED_TMX	
TONE_FAILED_READS	
INY_FAILED_READS	
FAILED_RCVRLK	

Table 5.2 Generic Rule Variables

possibilities are tested. The same functional limitation also applies to the detection of the DSCS III anomalies. An example of this situation is found in the TT&C #3 anomaly.

TT&C #3 is an anomaly which is common to both DSCS II and DSCS III. The following symptoms are indicative of a TT&C #3 anomaly: the satellite operator has no telemetry, and the RGF ground antenna operator reports the presence of a carrier and subcarrier signal, but an absence of telemetry modulation on the downlink signal. In this particular anomaly, the most probable cause of the anomaly is the satellite encoder. The standard procedure on DSCS II is to swap the encoder and if the anomaly is not resolved, swap the encrypters. If the anomaly is still unresolved, under DSCS II procedures, the operator would then swap the transmitter and, as a last resort, the encrypters would be bypassed. DSCS III does not follow these procedures. For a TT&C #3 anomaly, DSCS III procedures instruct the operator to swap the encoder. If the encoder swap does not resolve the anomaly, the operator then swaps the encoder power supply. In either case, GISMO swaps the encoder and assumes the recovery was successful. In real-world scenarios the expert system may need to swap multiple systems to resolve the anomaly, but as shown in this example, the paths taken to resolve the anomaly may be specific to the satellite constellation.

It is very difficult to quantify the success of the generic rule-base relative to its ability to detect all or part of the TT&C anomalies on both the DSCS II and DSCS III satellite constellations. Considering the fact that the initial detection is such a small part of the overall resolution of a satellite anomaly, the usefulness (success) of the generic rule-base (as an independent anomaly resolution capability) would seem to be limited. The standard scenarios implemented within the GISMO expert system display a limited initial anomaly recognition capability. After the initial generic detection has occurred, anomaly resolution and further diagnostic actions are taken through the use of satellite-specific rules and procedures. The results of the GISMO anomaly detection abilities may be

misleading since the complete anomaly recovery process and satellite-specific diagnostics are not included at this time. Once these items are added to the GISMO prototype, the usefulness of the generic concept may be more clearly determined. As a start, the simplifying assumptions used in the GISMO prototype will first need to be overcome.

5.2.2 Prototype Limitations to Overcome. The GISMO prototype is based on several simplifying assumptions. These assumptions were necessary to work around certain tool limitations and constraints. MAGIC must overcome these limitations to be a useful tool in the field of satellite operations. Below is a discussion of these assumptions and the problems that could result from using these assumptions in real-world situations. Some possible solutions to most of these assumptions are also discussed.

5.2.2.1 Multiple Point Failures. A multiple point failure exists when the symptoms of one anomaly are hidden by the symptoms of another anomaly and both anomalies become undetectable using normal methods. The diagnostic process of GISMO assumes no multiple point failures. Even though double failures do not occur very often, it is not very realistic to say it could never happen. This assumption also prevents the scenario in which a previously failed component could have some effect on the state of other components. To be an effective satellite controller, MAGIC must be designed to recognize and diagnose multiple point failures.

5.2.2.2 Communication Link Drop-Outs. In a satellite operations center, communication link drop-outs routinely occur several times a day. These drop-outs are also called data-hits. When a data-hit occurs during a satellite contact, the telemetry is momentarily disrupted and garbled, giving a false representation of the true satellite state. This is a common occurrence, so most operators simply ignore the

scrambled data. A computer, however, may have some trouble distinguishing between a data-hit and a satellite anomaly.

The MAGIC system would need to include a check for data-hits. For example, when a data-hit occurs, an abnormal number of anomalies will be flagged all at once. As a safety check for data-hit scenarios, the expert system could keep a count of the anomalies detected in a certain time span. If this number of detected anomalies grew too large too fast, then the expert system should flag a possible data-hit situation. At that time the expert system should wait for a few moments to allow the possible data-hit to clear and then begin its diagnosis again.

5.2.2.3 100% Recovery Success, Hard Failures, and Immediate Satellite Response. These three assumptions are grouped together because there is one solution that could remove all three assumptions. Before discussing the solution, each assumption and the real-world problems associated with the assumption will be discussed.

Because the GISMO prototype does not have a GUI satellite anomaly resolution system to interact with the operator and guide him or her through the recovery process, all simulated recovery procedures were assumed to be successful. Of course this is an unacceptable assumption for a real-world system. MAGIC will need a truly interactive recovery system, designed with procedures to handle contingency scenarios.

Another assumption appropriate for the anomalies addressed in the TT&C subsystem would not work in other subsystems. The assumption held that any component whose value fell below the lower red limit was a failed component. The component was treated as a hard failure with no ability to recover. This assumption would be invalid in the case of a battery voltage that fell below its lower red limit because it was being reconditioned. In this case the battery would be placed into a charge configuration and the battery current would then rise into the nominal limit range. In this scenario the

battery had not failed just because its current value fell below the lower red limit. The MAGIC system will need to address this issue.

Finally, the GISMO expert system assumes the response to a recovery procedure is transmitted from the satellite and received by the satellite control center instantaneously. Once GISMO detected an anomaly and simulated the recovery, it was necessary for the updated satellite state to appear in the very next set of satellite telemetry. A response from a satellite is not instantaneous; realistically, the response can take over two minutes. The expert system, meanwhile, has possibly analyzed over 120 sets of telemetry. Those 120 sets of telemetry reflect a failed state because the satellite response has not been received. The expert system must remember that the recovery procedure has been performed and that it must wait for a response to verify if the recovery was successful.

This memory capability is the solution to the last three assumptions discussed above. With the ability to remember what recovery procedures have been implemented, the expert system can perform further diagnosis if the initial recovery was unsuccessful. This eliminates the 100% successful recovery assumption. Also, with a memory capability, the expert system does not need to automatically assume a component has completely failed. The expert system can again perform the corrective actions and wait for the component sensor readings to move back into the nominal limit range. The following section discusses a possible approach to adding a memory capability to GISMO.

5.2.2.4 Memory. Variables within EXSYS could be used as a memory tool. A variable could be created for each recovery procedure and when a recovery procedure was invoked, the corresponding variable would be set to True to alert the system that this procedure has been used. Of course, each rule would need to be prefaced with a precondition, "if recovery procedure X has not been used." A general approach to implementing this suggestion is shown below.

IF Anomaly X is detected
and Recovery Procedure X has not been implemented
THEN Implement Recovery Procedure X

IF Anomaly X is detected
and Recovery Procedure X has been implemented
and Wait Time for Recovery Procedure X has been exceeded
THEN Implement Recovery Procedure Y

A more specific approach, using variables, is shown below.

IF anomaly is detected
and [Procedure_X_accomplished] = False
THEN Implement Recovery Procedure X
and [Procedure_X_implemented_at] IS GIVEN THE VALUE "Current Time"
and [Procedure_X_accomplished] IS GIVEN THE VALUE True

IF anomaly is detected
and [Procedure_X_accomplished] = True
and "Current Time" > [Procedure_X_implemented_at] + 2.2 minutes
THEN Implement Recovery Procedure Y
and [Procedure_Y_implemented_at] IS GIVEN THE VALUE "Current Time"
and [Procedure_Y_accomplished] IS GIVEN THE VALUE True

IF [Procedure_X_accomplished] = True
and "Current Time" < [Procedure_X_implemented_at] + 2.2 minutes
and satellite has responded as expected to recovery procedure X
THEN make updates to limit/status database
?? and [Procedure_X_accomplished] IS GIVEN THE VALUE False ??

Several issues arise from the implementation of the above described memory tool. The first issue to be resolved is: When should the updates to the limits/status database be made? Another item of concern relates to the required wait time for the satellite recovery response and whether the procedure variable should be reset within the satellite contact in which it was used.

The issue of when to update the satellite's limits/status database after a recovery procedure is implemented needs to be resolved. The options to resolve this issue are: (1) wait until the satellite state verifies a successful recovery procedure to make the updates, or (2) make the updates to the limits/status database immediately after the recovery procedure is performed and before the telemetry has changed.

My recommendation would be to wait until the recovery is complete before making the database updates. This is because, for example, in the TT&C #3 and #4 passplans components are swapped without knowing definitely if they are the cause of the problem. If, after swapping the first most probable cause of the problem, the problem is not resolved then that unit is probably not failed unless this is a multiple point failure scenario. The swapped component still has a viable redundancy and its redundancy variable should not be changed to NO. It is probably better, therefore, to wait until the anomaly has been completely resolved before updates to the limits/status database are made.

There are two possibilities to resolving the necessary wait time for a particular recovery response. The MAGIC system will have knowledge of where the telemetry point of interest is located within the telemetry stream. With this knowledge and the knowledge of where the system is within the stream at that time, it would be a simple task to calculate the remaining time necessary to wait for the telemetry update from the satellite. Another option would be to default to the worst-case scenario and set all wait times equal to the longest wait time for any telemetry point. This is equivalent to the time necessary for the satellite to transmit one master frame (every telemetry point). The time to transmit one master frame for a DSCS II satellite is approximately 131 seconds. DSCS III takes only 61.44 seconds to transmit a master frame of telemetry because the satellite transmits its telemetry at 1000 bps, whereas DSCS II transmits at a rate of 250 bps.

The final issue pertains to the last post-condition of the last rule shown in the example above. Should the "Procedure_X_accomplished" variable be reset to False after the recovery is complete? This would allow the procedure to be used again within the same satellite contact if necessary; otherwise, the procedure would be locked out and unusable until the next satellite contact. If the variable is reset to False the expert system must be in forward-chaining mode or EXSYS will backward chain over the rule to find the full and final value of the variable if needed. Rule ordering would be important and the last rule in the above example would need to be placed at the end of the set of rules that use the "Procedure_X_accomplished" variable. It is not very likely for the same type of anomaly to occur twice within the same support, but it is not impossible either.

5.2.3 Concerns and Opinions. As a result of these simplifying assumptions, necessary to develop the GISMO prototype, there are some concerns on the validity of the results. Also, throughout this thesis effort I have formulated some opinions on the future of generic satellite controllers. Both my concern and opinions are discussed below.

5.2.3.1 Concerns. The first concern is that the results may be limited due to the implementation of only one satellite subsystem. For example, when a component falls below its lower red limit and can later rise back into its nominal, green range was not an issue in the TT&C subsystem. This was because, once the component had exceeded a particular limit, it was immediately considered failed and was no longer expected to operate again. This is different from the Electrical Power Subsystem (EPS), where batteries may fall below their lower red limit during reconditioning and later climb back into their nominal range.

Also, I was originally concerned that the results of this thesis may have been limited due to the satellites chosen, because both were geosynchronous communications

satellites. These concerns have now been addressed since the Navstar Global Positioning System (GPS) satellite was added to the prototype. The results of this merger is discussed below in Section 5.3.

The validity of the results is also a concern because the prototype was designed to diagnose anomalies on those TT&C systems designed to operate on the AFSCN. Those satellites which utilized only dedicated ground antennas were not addressed, because the engineers of those TT&C subsystems had no restrictions on the design and more specific designs evolved in these systems. Future efforts may want to add such systems, even though MAGIC will be successful if it can only operate on satellites designed for the AFSCN.

5.2.3.2 Opinions. At the beginning of this thesis effort I chose a domain for my expert system. I chose the TT&C subsystem of the DSCS II and DSCS III satellites based on its simplicity and generality. The TT&C subsystem is somewhat standard on most satellites designed to operate on the AFSCN. I feel if the results of this prototype, designed on such a limited domain, is unsatisfactory, then there is little hope for acceptable success on other satellite subsystems.

Also, as satellites are added to the prototype, I feel the generic rule-base will be reduced in size. This is because the satellite being added may have a particular design that prevents it from fitting into a particular rule; therefore, the rule must be removed and made into specific rules. If a generic rule does not apply to all satellites then it is not generic.

A possible solution to this problem would be to add another level of abstraction to the prototype. More specifically, a precondition could be added to each rule that states "if the satellite of concern has this particular component." This would add more overhead to the system, but at the same time, reduce the need to delete a generic rule if it does not

apply to all satellites. This would require more definitions to be available in the satellite-specific databases to define what components are and are not available on the particular satellite. The implementation of this idea would prevent the generic rule-base from reducing in size if a satellite's design did not include a particular component tested within the generic rule-base.

The major factor causing difficulty in designing generic rules is the non-standardization between different satellite constellations. If the results of this prototype are found to be unacceptable, a possible option may be to address the design of a MAGIC system for satellites with the same mission. Another possibility might be a MAGIC system designed for satellites built by one particular manufacturer.

Finally, I feel the best hope for the success of a generic satellite controller lies in the success of transforming multiple satellite telemetry streams into a generic telemetry stream the expert system will operate on. This capability would be part of the telemetry preprocessing system. This could mean consolidating multiple telemetry points into one point that can convey the same information. This, in itself, could require some rule-base reasoning to implement. This would not eliminate the need to check the status of the additional points, but it would allow the generic rules to diagnose an anomaly using the generic telemetry point derived from several specific points.

5.2.4 Items to Investigate.

5.2.4.1 More Powerful Software Tools. To overcome certain limitations of the current prototype (described in Section 4.3), Phillips Laboratory has purchased some new software that will need to be incorporated into the current prototype.

Phillips Laboratory purchased the Linkables software package, developed by EXSYS, and the Microsoft NT operating system. The EXSYS Linkables allows the

EXSYS expert system to embed subroutines within the system. The expert system rules can make calls to these embedded subroutines to satisfy procedural needs. Also, all internal variable values are global, allowing the embedded subroutines access to all of the variable values held within EXSYS. The NT operating system allows true multitasking. Using NT, the generic expert system, after detecting an anomaly, can make a subroutine call to the proper recovery procedure and continue running its rules while the recovery procedure works with the operator to resolve the detected anomaly.

The use of embedded subroutines will aid in the development of generic rules. This is because satellite-specific recovery procedures and required variable updates, originally placed directly in the consequence of the rule, can be replaced with a simple subroutine call. This subroutine call can be made using an embedded variable which makes the call generic because the rule can call multiple recovery procedures using one subroutine call. An example of this is shown below.

```
IF      Preconditions are met
THEN    Call Recovery Subroutine [[Constellation]].rcv
```

In this example, if the preconditions are met, EXSYS will call the recovery subroutine. The name of the subroutine will be the value held by the Constellation variable with a .rcv extension. If, for a particular satellite contact, the variable Constellation holds the value DSCSII and the preconditions in the rule above are met, EXSYS will call the subroutine named DSCSII.rcv.

The addition of these new software packages, along with the capabilities of the Telemetry Analysis System (TAS) will overcome the current work-arounds used in the GISMO prototype.

5.2.4.2 Telemetry Preprocessing. As stated above, if more extensive telemetry preprocessing was accomplished, more commonalities could be formed.

Table 5.3 shows some points that could be combined into one variable for use in the generic rule-base to diagnose common anomalies on both satellites.

<u>DSCS II</u>	<u>DSCS III</u>	<u>GENERIC</u>
TMXTP.....	TMXAP.....	Active_TMX_Power
	TMXBP.....	
TMX_T.....	TMXAT.....	Active_TMX_Temp
	TMXBT.....	
CMD1CK.....	CMDVER.....	CMDCK
CMD2CK.....		
ENCRP.....	K46A28.....	ENCRP
	K46B28.....	

Table 5.3 Combined Telemetry Points

The DSCS II telemetry point, TMXTP, represents the output power of the selected transmitter while DSCS III has two separate telemetry points for each transmitter output power. Knowing which transmitter is selected, through the transmitter status variable, a telemetry preprocessor could take the value of the telemetry point corresponding to the active transmitter and assign that value to the new generic variable. The same process would work for the DSCS II and DSCS III transmitter temperature telemetry. Also, the preprocessor could combine the command processor command check telemetry (CMD_CK/CMDVER) into one telemetry point representing the active command processor. The last point shown above is the encrypter status. The DSCS II has a discrete status telemetry point to represent the state of its encrypter while DSCS III has a converter status for each of its encrypters. The preprocessor could determine from

the value of the converters and the expected status in the database, which encrypter was selected and place that information in a generic variable.

With the telemetry preprocessing shown above, the results of Table 5.1 have changed. The new commonality results are shown in Table 5.4.

	Total TT&C Tlm Points (X)	# Common with other satellite (Y)	Percent in common (Y/X)
DSCS II	45	22	48.89%
DSCS III	66	22	33.33%

Table 5.4 Commonalities After Telemetry Preprocessing

When several specific telemetry points are consolidated into one generic variable, some information is lost. To insure the information is not lost, the expert system would load all telemetry points and do some fundamental checks on each individual point, to include the basic limit checks and using the values for trending purposes. Once this is done the generic variable, which may be a consolidation of several telemetry points, could be used by the generic rule-base to diagnose a component failure.

5.2.4.3 Additional Suggestions. To provide a more complete evaluation, another satellite subsystem should be added to the prototype. I believe the Electrical Power Subsystem (EPS) would be a good choice. There are some notable similarities in the use of batteries on different satellites, but the operation of those batteries and power supplies could vary. The results of a second satellite subsystem combined with the results

of this research should give a clear picture of the possibility of success for the generic concept.

If the rule structure developed in this research does not provide desirable results for the sponsor, maybe another knowledge representation should be attempted. The use of a higher level of abstraction, such as placing the necessary satellite-specific checks in a look-up table, could be useful in developing a more generic rule structure.

If all suggested options have been exhausted without acceptable results, I would suggest the use of additional reasoning methodologies. A possible combination of model-based and case-based reasoning with the existing rule-base could provide some powerful capabilities to the generic intelligent controller.

5.3 Initial GISMO Prototype with a Third Satellite Extension

With the prototype in place and the generic rules structured to take advantage of those commonalities between DSCS II and DSCS III, a third satellite was analyzed for compatibility with the existing GISMO prototype. This was done to test the viability of the GISMO expert system and the generic concept. The added satellite was the Navstar Global Positioning System (GPS) satellite, maintained and operated by the First Satellite Operations Squadron (1SOPS) at Falcon AFB in Colorado Springs, Colorado. GPS is a space-based navigation network that provides accurate position, velocity, and time to its users worldwide. Since there was concern that the DSCS II and DSCS III satellites were too similar to justify a valid test, GPS was selected to address these concerns. GPS is a three-axis stabilized, semi-synchronous satellite. Semi-synchronous satellites are located in an orbit that is one half the distance of a geosynchronous satellite from earth. A detailed description of the GPS TT&C subsystem can be found in Section A.5.

An in-depth analysis of the existing rules and the GPS TT&C subsystem failures was performed to see if the generic rules could correctly diagnose anomalies on GPS as

well as DSCS II and DSCS III. For this analysis, a few assumptions were made and are listed below.

- Only GPS Block IIA satellite models are considered
- Nominal configuration is Automatic Turn-On (ATO) mode, although I did consider the usefulness of the rules for satellites in COM and CTO mode
- No PRN ranging data is collected. PRN is inhibited on both receivers.

During the study of the functionality of the GPS TT&C subsystem, several distinct differences and similarities between GPS and the two DSCS satellites became apparent.

5.3.1 TT&C Differences. There are several major differences between the functionality of the GPS TT&C subsystem and that of the DSCS satellites. These differences caused some difficulty in merging GPS into the existing GISMO rule structure.

5.3.1.1 Operating Modes. GPS has three possible modes of operation. The nominal mode is the Automatic Turn-On mode (ATO). In ATO mode the satellite telemetry transmitter is automatically turned on at the presence of a signal generated within receiver #1 in response to a modulated uplink signal. The transmitter is automatically turned off sixteen seconds after the loss of that signal. Another operating mode is the Continuous On Mode (COM). Satellites that have experienced some problems or failures in their TT&C subsystem may be placed into this configuration. In the COM mode the transmitter is always on, just like the operation of the transmitter on the DSCS satellites. The last mode is the Command Turn-On (CTO) mode. This mode is used if other GPS satellites are in close vicinity and possible Radio Frequency Interference (RFI) could occur causing random transmitter turn-on in ATO mode, or downlink interference in COM mode. In the CTO mode, the satellite operator must send the transmitter-on command to receive telemetry.

5.3.1.2 Receiver Operation. Unlike the paired DSCS receivers that are designed to search for two different frequencies, the paired GPS receivers search for the same frequency. This means both GPS receivers will lock onto the uplinked signal instead of the single receiver lock-up on the DSCS satellites. Since both GPS receivers lock onto the uplinked signal, both receivers process the uplinked signal and strip off the command tones and send the 1, 0, and S-tones (discussed in Section A.5.1) to their corresponding decrypter (INY). Only the INY which has the decryption scheme to match the uplinked encryption scheme will decrypt the command.

Both receivers also strip off the PRN signal and send it to the active transmitter. This causes interference with the downlink signal; therefore, one of the receivers must be inhibited. If a receiver is inhibited it cannot send the PRN data to the transmitter.

The 1SOPS does not collect PRN ranging data for GPS because the addition of PRN on the downlink weakens the signal enough to cause difficulty for the GPS dedicated Ground Antennas (GA). Because the mission of GPS is one of precision position calculations, the GPS satellite knows where it is and it is not necessary to collect ranging data; therefore, PRN is inhibited on both GPS receivers.

5.3.1.3 Cross-Strap Timer. GPS has a cross-strap timer that will automatically cross-strap the receivers and KIR decrypters if no commands have been processed by the satellite within six days. This is a built-in safety feature which can be disabled by ground command. Each time a command is processed by the Dual Command Decoder (DCD), the timer is reset. If there is a component failure on the uplink side of the TT&C subsystem, procedures require the cross-strap timer to be disabled. This is done to prevent the satellite from cross-strapping itself into a failed configuration in which commanding is then made impossible.

5.3.1.4 Encrypter Cross-Strap. GPS does not have the capability to bypass its downlink encrypters as the DSCS II and DSCS III satellites do. If both encrypters fail on the GPS satellite, no telemetry will ever be transmitted by the satellite again. Though it is not a preferred configuration, if both encrypters were failed, both DSCS II and DSCS III could still provide satellite telemetry by bypassing the downlink encrypters and sending the telemetry in clear text mode.

5.3.2 TT&C Similarities. Even with the differences described above taken into consideration, the similarities of the TT&C functions are still encouraging. For instance, though there are certain differences in the function of the GPS receiver, it maintains a basic function identical to the one performed by the receivers of both DSCS satellites. The GPS receiver strips the command data off of the uplinked carrier, and like the DSCS III receiver, it outputs the 1, 0, and S pulses to the front end of the command processing unit (DCD).

Another similarity is the function of the Decoder/Activate (D/A) signal. The GPS decoder activate (D/A) or tones detector is located within the receiver similar to the DSCS III AMSYNC detector. The function of the D/A component is exactly the same as that of the DSCS II and DSCS III satellites. The D/A component searches for tones on the uplinked signal and becomes active at the presence of tones.

5.3.2.1 GPS Telemetry Commonalities. Once the basic functionality of the GPS TT&C subsystem was understood, each telemetry point was analyzed and compared to the generic points extracted from the DSCS satellites. Table 5.5 shows the outcome of the telemetry analysis.

	Total TT&C Tlm Points (X)	# Common with other satellites (Y)	Percent in common (Y/X)
DSCS II	45	12	26.67%
DSCS III	66	12	18.18%
GPS	53	12	22.64%

Table 5.5 GPS Commonalities with DSCS II and DSCS III

Surprisingly, though, all but one of the telemetry points used by the generic rule-base were found to be common with the GPS TT&C subsystem.

5.3.3 Merging Difficulties. For the GISMO expert system to operate correctly on the GPS telemetry, several changes would be necessary. Below is a discussion of some items which may cause difficulty when operating on GPS telemetry and also some rule-specific information on how some rules worked and others did not.

The first item noticed, which could cause some difficulty, was the lack of an active/passive status point on the GPS telemetry frames. DSCS II and DSCS III have an ACTIVE/PASSIVE status displayed on one of their telemetry frames. This value is derived on the ground and is, therefore, not transmitted by the satellite. The ACTIVE/PASSIVE status represents the radiating status of the ground antenna. Currently, the ACTIVE/PASSIVE status is not available on the GPS telemetry consoles, but I assume this would be a trivial thing to accomplish. Possibly some internal procedure could continuously look for the active directive to be sent and when it is sent the procedure would set the AACTV_PASV variable to "ACTIVE" within EXSYS.

In several of the GISMO rules, the pre-condition, "Good Range is/is not Received" is used to diagnose transmitter failures. Since range is not collected on GPS, these rules would not work properly on the GPS telemetry. To get around this, another precondition could be added to these rules stating "Range is/is not collected". This information would need to be defined in some satellite-specific database. I feel this would only be necessary for GPS because I believe all other satellites do collect range data.

The next problem encountered was with the generic angles-match rule. This rule checks to see if the expected azimuth and elevation angles, found in the ephemeris database, match the actual ground antenna angles. Currently, the simplified ephemeris database used by the GISMO expert system contains one set of azimuth (AZ) and elevation (EL) angles per satellite per Remote Ground Facility (RGF). To accommodate a satellite which moves with respect to the earth, the prototype ephemeris database would need to expand to provide AZ and EL angles for each satellite at each visible RGF accurate to the minute of the day. This increases the size of the database but is not a difficult task. GISMO would need to search the database, not only on the current RGF in use, but also on the current time.

Another difficulty encountered in merging GPS was the set of generic INY failure rules. The generic INY failure rules do not work for GPS because GPS does not have a sensor point for INY status. The only way to tell if the KIRs (INYs) are operational is to check for the presence of the Vehicle Command Count (VCC) that is telemetered by the KIR decrypters. Instead of simply checking the value of the INY status telemetry point to diagnose a failed INY, GPS must check the following (in ATO mode):

- Telemetry Present?
- Receiver lock nominal?
- Tones nominal?
- DCD nominal?
- VCC nominal?

If the VCC value is bad, the GPS approved procedure does not automatically configure for the alternate decrypter, but instead cycles the uplink modulation twice. Afterwards, if no valid VCC exists, then it configures for the alternate decrypter and sets the redundancy value to NO. Also, to use the alternate command path, the alternate DCD must be commanded on which is not a requirement on the DSCS satellites. The above method of diagnosing a failed INY is not the same as that used in the current GISMO design. In GISMO, the INY bi-level status is checked directly to determine the status of the DSCS II or DSCS III INY. This diagnostic process would not work for GPS.

Of course, some GPS specific rules could be placed above the generic INY rules which would test the VCC values. If the VCC values were abnormal, a GPS rule could recommend the cycling of uplink modulation as an initial attempt to fix the problem. If, after cycling the uplink modulation twice, no recovery occurred, then the GPS specific rule could set the INY qualifier to abnormal and let the generic rule handle the rest.

Another problem was found in the way the GPS satellite telemetry represents its primary and secondary components. Within the GISMO design, the "Primary" and "Secondary" variable definition tactic helped to consolidate some satellite-specific rules into generic ones. GPS is not as consistent as DSCS II and DSCS III in its component status representation. GPS uses the numbers 1 and 2 and the letters A and B to represent its primary and secondary units. This may require that the generic rule-base replace the generic variable "Primary" with a generic variable specific to the failure being diagnosed. For example, in a transmitter failure rule the "Primary" variable could be replaced with "Primary_tmx". This variable would be defined in the satellite-specific database and would represent the value of that satellite's primary transmitter, whether it be ONE, or 1, or A. Of course, this would mean that there would need to be a definition variable for every component used in the generic rule-base.

5.3.4 Results and Findings. Every generic rule within the GISMO expert system was analyzed to determine if the rule would operate correctly on GPS telemetry. During the analysis, I considered all three possible GPS operating modes while performing the rule analysis. Table 5.6 shows the outcome of this rule compatibility analysis.

Generic Rules	Operational Modes		
	ATO	COM	CTO
Bad Downlink Configuration	E	E	E
Bad Uplink Configuration	E	A	E
Range Check	C	C	C
Bad Range	F	F	F
Failed Tone Detector	G	A	G
Failed Tone w/no redundancy	n/a	A	n/a
Failed INY	D	D	D
Failed INY w/no redundancy	n/a	B	n/a
Failed Receiver	G	A	G
Failed Receiver w/no redundancy	C	A	C
Failed Subcarrier Generator	A	A	A
Failed Transmitter	B	B	C
Failed Transmitter w/no redund.	B	B	C
Failed Primary Transmitter	B	B	C
Failed Secondary Transmitter	B	B	C
Bad Tlm Synch (failed Encoder)	A	A	A
Failed Encoder	F	F	F
Failed Encoder w/no redundancy	B	B	B
Failed Primary Encoder	B	B	B
Failed Secondary Encoder	B	B	B

LEGEND:

A - Rule works as is

B - Rule will work with a few simple alterations

C - Rule will not work at all

D - Rule does not apply to GPS and will need to be made into a specific rule

E - Rule does not apply to GPS but its existence does not effect the results of the run

F - Rule does not apply to GPS, but a precondition could be added to the rule so that the rule does not fire for GPS and will not need to be replaced with satellite specific rules.

G - Rule will work in certain failure scenarios

n/a - Not-applicable because the satellite would not be in this mode if in this configuration

Table 5.6 Generic Rule Compatibility with GPS

As I analyzed each rule within the GISMO prototype to see if the rule could operate correctly on GPS telemetry, I noticed two items of interest. The first interesting finding was the redundancy of preconditions in certain rules. Certain generic rules that effectively detected anomalies on a GPS satellite in COM mode could also detect the same anomaly on a satellite in ATO mode, but some of the pre-conditions of the rule would be redundant. For example:

```
IF    Telemetry is present
    and [AACTV_PASV] = "ACTIVE"
    and [SPE] = "E"
THEN Do something
```

The above preconditions are necessary for a GPS satellite in COM mode, but if the satellite is configured for ATO mode, some of the above preconditions are redundant. If a GPS satellite is configured for ATO (nominal) mode and telemetry is present, then the ground antenna must be active and S-pulses must be enabled (SPE = E). In this nominal configuration, the only precondition required in the above example would be the "Telemetry is present" precondition. It does not make the diagnosis incorrect to perform these additional checks.

The second item of interest pertains to the compatibility of DSCS II and DSCS III satellite-specific rules with GPS. During my analysis, I found several instances in which a DSCS II specific rule could work for GPS, but not DSCS III. In other cases, I found situations where a DSCS III specific rule could work for GPS telemetry, but not for DSCS II. For example:

- The DSCS II Tones #1 and #2 failure would work for GPS, but not DSCS III
- The DSCS II INY #1 and #2 failure would work for GPS, but not DSCS III
- The DSCS III Receiver Failure rules would work for GPS (if telemetry exists), but not DSCS II

This shows support for the hypothesis that there are commonalities between satellites, yet some subtle differences cause the separation of the diagnostic rules.

The result of adding GPS to the GISMO structure was a partial success. With a few adjustments, all generic rules, except one, in the GISMO expert system worked correctly on the GPS telemetry as long as the satellite was in COM mode. COM mode maintains a continuous downlink telemetry stream like the DSCS II and DSCS III satellites. The problems occurred when the satellite was configured for CTO or ATO

mode (the nominal operating mode of GPS). Only a few of the generic rules would handle a satellite in ATO mode. The outcome of this test gives good support for the success of adding other satellites that maintain a continuous downlink telemetry stream, but it does not give much support for those satellites that turn off their telemetry transmitter at the termination of each satellite contact.

The merger of the GPS satellite anomaly detection requirements was quite successful if you assert the simplifying assumption that all GPS satellites are continuously configured for COM mode, otherwise the merger resulted in limited success.

5.4 Summary

This chapter has detailed the results of the initial two satellite model and the results of adding the GPS satellite to the existing GISMO prototype. Several limitations of GISMO were described along with recommended techniques to overcome the limitations. Some of my personal opinions and concerns about the validity of the results were also described. Last, the proposed benefits of extensive telemetry preprocessing, along with a specific example, was described. The following chapter will summarize the entire thesis effort.

VI. Summary and Conclusions

The purpose of satellite operations, within the military, is to operate and maintain all DOD satellites. These priority systems are monitored daily to ensure these satellites are maintained in the best possible condition. The job of satellite operations is very tedious and some aspects are repetitive. The need for some automated assistance in the area of satellite operations has never been as prevalent as it is today.

Today, there exists the generic satellite operator who receives little to no training on the specifics of any satellite, yet is expected to command and control these multi-million dollar assets. To provide assistance to these generic operators, Phillips Laboratory has developed the concept of a Multimission Advanced Ground Intelligent Control (MAGIC) system. This system would assist the satellite operator in all aspects of his or her job. The MAGIC system is cost efficient because its conceptual design includes the ability to command and control multiple satellite constellations. This is in contrast with current R&D in which intelligent controllers are designed to control one specific constellation of satellites.

This research explored existing questions and generated new questions on the implementation of a generic rule-base to detect satellite anomalies on satellites of multiple constellations. To scope the problem, two satellites were chosen (DSCS II and DSCS III) to incorporate into the knowledge base. The problem was further scoped to the Telemetry, Tracking, and Commanding (TT&C) subsystem of the DSCS II and DSCS III satellites. An analysis of the two subsystems was performed to extract any commonalities. Afterwards, the rule-base was developed to detect known anomalies on the two TT&C subsystems using the extracted commonalities.

A certain rule structure evolved through the process of developing the prototype. The prototype included a generic and satellite-specific rule-base. The satellite-specific

rule-base was used to test preconditions specific to a particular satellite. The basic flow of anomaly detection in the prototype included a generic rule testing common preconditions and if the common preconditions were satisfied, the generic rule would set an internal variable as the consequence of the generic rule. This internal variable would flag a satellite-specific rule to check some satellite-specific preconditions and if they were also successful, the anomaly would be diagnosed and the recovery procedure would be simulated. There were very few cases in which a generic rule or group of generic rules could completely diagnose the anomaly without the assistance of some satellite-specific rule(s).

The results of the two-satellite prototype provided support for the feasibility of a generic satellite intelligent controller. To add further support for the generic rule-base proof of concept, a third satellite (GPS) was added to the original prototype. When the GPS satellite was configured for Continuous-On-Mode (COM), the results of the third satellite extension were quite successful, yet limited success was obtained when the GPS satellite was configured for Automatic-Turn-On (ATO) or Command-Turn-On (CTO) modes.

I feel that there were enough commonalities extracted and enough anomalies detected to justify further research. The results of this thesis provides support for the generic rule-base concept, but much more work needs to be done before the final decision on the feasibility of the generic satellite intelligent controller is made.

The current prototype limitations need to be overcome. A memory technique needs to be added to the prototype to allow several levels of diagnosis. With this capability it will become more clear, how much of the diagnostic process is generic and how much is specific. Other limitations to overcome include:

- The inability to recognize a communication drop-out
- The inability to handle deviations from the planned anomaly recovery procedure

- The inability to allow for a delay in satellite response to a recovery procedure
- The inability to treat an extreme telemetry value as a soft rather than hard failure

I do not feel the generic "Proof of Concept" was completely successful, but I would say a small limited generic rule-base is very possible with the techniques used in this thesis. If, after the recommended changes discussed above are successfully incorporated into the prototype, the capabilities are not satisfactory, maybe other approaches should be attempted.

One approach might be to limit the domain of the intelligent controller to satellite constellations which belong to a common mission. This thesis showed good results between the DSCS II and DSCS III satellites which share the same mission of providing worldwide military communications. There may be some inherent commonalities between satellites of common missions. Another approach might be to limit the domain to satellites designed by a common builder. This was not a contributing factor in the results of this particular research, but it may hold true that satellite design engineers incorporate certain commonalities into every satellite they design.

If the above suggestions do not produce acceptable results, it may be beneficial to re-evaluate the tools being used in the current prototype. It may be true that the domain of satellite operations is too dynamic for a rule-based system to control by itself. In addition to rule-base reasoning, maybe the capabilities of model-based and case-based reasoning should be added to overcome the weaknesses of a strictly rule-based system. Model-based and case-based reasoning have capabilities which could be very useful. I believe a lot of power could result from a mixture of reasoning methodologies similar to the ideas of Captain Jim Skinner discussed in Chapter II. Pulling on the strengths of several reasoning methodologies to overcome the individual weaknesses may have much success in the field of generic intelligent satellite controllers.

APPENDIX A. The TT&C Subsystem

This appendix details the operation of the Telemetry, Tracking and Commanding (TT&C) subsystem found on the DSCS II, DSCS III and GPS satellites.

A.1 DSCS II TT&C Uplink Section

The Defense Satellite Communications System Phase II (DSCS II) is a second generation military communications satellite system developed to replace the Initial DSCS (IDSCS) program. The DSCS II provides communications for the Worldwide Military Command and Control System (WWMCCS) to include the National Command Authority (NCA), Commander in Chief's (CINCs) and designated military command elements. (20:1)

The DSCS II satellites are placed in a geosynchronous orbit, positioned over the equator, at approximately 22,000 nautical miles above the earth at key longitudinal points. At this altitude, the satellites maintain a 24 hour rotation period. This allows the satellites to maintain a constant position with respect to the earth.

The DSCS II satellite was built by TRW. The DSCS II satellite was designed to operate for approximately five years, but most have performed operations much longer than the intended five years. This section along with Section A.3 explain the functionality of the DSCS II TT&C subsystem.

Figure A.1 shows a simplified block diagram of the DSCS II TT&C uplink section. Beginning at the left, the omnidirectional (OMNI) antenna receives and transmits signals in all directions. The diplexer allows the OMNI antenna to receive and transmit simultaneously and the hybrid splits the uplinked signal power equally between the two receivers. Both receivers are always powered and sweeping several KHz on either side of their particular channel. Receiver one sweeps about Space to Ground Link System




Figure A.1 Uplink Section of the DSCS II TT&C Subsystem (19:2)

(SGLS) channel 15 while receiver two sweeps in search of SGLS channel 16. When either receiver locks onto an uplinked carrier signal it sends a "Signal Present" signal to the Electrical Integration Assembly (EIA) so that the EIA will power-up the command decrypters. Following the receivers are the signal conditioners (SIGCON). The sigcons are always on searching for the presence of tones on the uplink signal. Tones are frequencies on the uplink used to represent a 1, 0, or S-pulse. Different combinations of 1's and 0's are used to represent a satellite command, with the S-pulses used before and after each command to help the satellite distinguish between the end of one command and the beginning of another. Both the receiver and signal conditioner pair are continuously powered by a common receiver converter. The KI-23 command decrypters, or more commonly known as INYs, decrypt the uplinked command according to the command key used on the ground to encrypt the command. For example, if the secondary command key

is used to encrypt the command before it is transmitted to the satellite, then INY two, on the satellite, is used to decrypt the command. This is made possible because there exists a cross-strap capability between the signal conditioners and the INYs. This cross-strap capability allows any signal uplinked to the satellite to reach both INYs automatically. The Data Interface Assembly (DIA) serves as a buffer and filter between the INYs and the Electrical Integration Assembly (EIA). The EIA processes the command after it has been decrypted and ensures the proper subsystem executes the command. Any command designated for the despun platform is sent to the Switching Logic Assembly (SLA) for execution. (19:2,3)

A typical uplink scenario would occur as follows. The ground operator would send a command consisting of a combination of 63 1's and 0's across the communications link to the designated ground antenna. Once there, the command, along with the Pseudo Random Noise (PRN) ranging data, is modulated onto the requested uplink channel and transmitted to the satellite. Before the satellite contact begins, the satellite operator briefs the ground antenna operator which uplink channel will be used for this particular support. This would be the channel used to transmit commands and ranging data to the satellite. The satellite's omni antenna would collect the signal and route it to the diplexer which would split the power of the signal and route half of the power to receiver one and the other half to receiver two. If channel 15 was used as the uplink channel, receiver one would lock onto the signal and send out a "Signal Present" signal to the EIA which in turn sends a signal through the DIA to the INYs to power up both INYs. The receiver strips the Pseudo Random Noise (PRN) ranging signal from the uplink carrier and sends it directly to the selected Dual Baseband Assembly Unit (DBAU) of the TT&C downlink section. The DBAU modulates the PRN onto the subcarrier and sends it to the transmitter to be transmitted down to the ground antenna where it is used to calculate the range of the satellite. The receiver also strips the 63 bits of command data from the uplink signal and

sends it to both sigcons. The sigcons are continuously powered by the same converter that powers the corresponding receivers. Telemetry is misleading for the signal conditioner because it reads NONE when no uplink tones are present as if the sigcon was not powered, but it is continuously powered and searching for tones. Both sigcons convert the uplinked command to digital data and send it, along with a clock pulse, to their corresponding INY. Only the INY that holds the decryption scheme that matches the one used to encrypt the command on the ground, will decrypt the command. The other INY will just ignore the data sent to it. The proper INY compares a Ground Command Count (GCC), which is part of the 63 bit uplinked command, to the Vehicle Command Count (VCC) stored within the INY. If the two numbers match, the command is considered authenticated and is then decrypted and sent through the DIA and into the EIA for command processing.

It is important to note where the cross-strapping exist. All uplink signals are sent to both signal conditioners due to the cross-strapping between the receivers and the signal conditioners. Both signal conditioners send the data to their corresponding INY. Only the INY with the matching decryption scheme will authenticate and decrypt the command. The DIA and EIA are internally redundant but not cross-strapable. The DIA/EIA pair is selected based on what INY is used. There is cross-strapping between the EIA and the SLA, allowing either EIA component to send commands to either SLA component.

A.2 DSCS III TT&C Uplink Section

The Defense Satellite Communications System Phase III (DSCS III) is a third generation military communications satellite system developed to replace the DSCS II program. The DSCS III provides secure, anti-jam communications for the Worldwide Military Command and Control System (WWMCCS) to include the National Command Authority (NCA), Commander in Chief's (CINCs) and designated military command

elements. A secondary payload provides a back-up function to transmit the Emergency Action Message (EAM) and Single Integration Operation Plan (SIOP) in case of a national emergency.

Just like DSCS II, the DSCS III satellites are placed in a geosynchronous orbit, positioned over the equator, at approximately 22,000 nautical miles above the earth at key longitudinal points. At this altitude, the satellites maintain a 24 hour rotation period allowing the satellites to maintain a constant position with respect to the earth.

The DSCS III satellite was built by GE and designed to perform operations for ten years. DSCS III is different from the average satellite in that it has essentially two TT&C subsystems. The one used and discussed in this research is the S-band TT&C subsystem designed for use with the AFSCN. The other TT&C subsystem, on the DSCS III satellite, operates in the X-band area of the electro-magnetic spectrum. It was designed to operate with dedicated ground antennas. DSCS III has dedicated ground antennas to monitor and control the mission payload 24 hours a day. The X-band TT&C subsystem was designed for the exclusive use of these dedicated antenna sites. The X and S-band sections have some independent and some shared components. (21:1)

TT&C subsystems designed for dedicated ground antennas may not follow the basic functionality, or contain the commonalities, found in those designed for the AFSCN. This thesis is focused on those satellite TT&C subsystems designed to operate on the AFSCN. Future efforts may want to expand and add those dedicated systems such as the DSCS III X-band system.

This section along with Section A.4 explain the functionality of the DSCS III TT&C subsystem. Figure A.2 shows a simplified block diagram of the DSCS III S-band TT&C uplink section.




Figure A.2 Uplink Section of the DSCS III TT&C Subsystem (21:9)

Starting at the left of figure A.2, the +Z antenna collects the uplinked S-band signal. The -Z antenna is located on the opposite face of the satellite and is used if the satellite is in an anomalous attitude. The Hybrid and Diplexers operate the same as on the DSCS II satellite. Both receivers receive the uplinked signal, but only the one corresponding to the uplinked channel will lock onto the signal. Receiver A continuously sweeps a certain range of frequencies centered on SGLS channel 12 and receiver B sweeps about channel 16. The function performed by the signal conditioner on the DSCS II satellite is accomplished within the receiver of the DSCS III vehicle. The receiver strips the PRN ranging data off of the uplink and sends it to the selected transmitter. It also strips the command tones and converts them into 1, 0, and S pulses for the decrypters.

The selected KIR-23 (INY) receives the command data from the receiver and authenticates it in the same manner described for DSCS II. Only one INY receives the signal on the DSCS III because the cross-strap link, between the receivers and the INYs, is reconfigured by ground command only. Once authenticated, the command is decrypted and sent to the Command Decoder (CD A/B). The CD processes the command and ensures the command is executed properly, similar to the function performed by the DSCS II EIA component. (21:8)

A typical uplink signal on channel 12 would be accepted by receiver A. Receiver A would strip the PRN ranging code and send it to the active transmitter. When tones were detected on the uplink, the receiver would send out an AM Sync signal, analogous to the Signal Present signal used on DSCS II, to the CD. This signal triggers the South Panel Power Controller (SPPC) to power the INYs. Unlike DSCS II, there is a cross-strap capability just before the INYs. This cross-strap must be commanded to switch from one INY to another. Now, with the INYs powered, an uplinked command will travel from the locked receiver to the appropriate INY. After authentication, the command will be sent to the selected CD. The INYs are cross-strapped to the CDs and do not require commanding to connect to either CD. At the CD, the command is processed and then executed by the appropriate subsystem.

Due to the function of the antenna diplexer, the uplinked commands are processed and executed at the same time as the downlink section of the TT&C subsystem is transmitting satellite sensor data (telemetry). The following sections describe the DSCS II and DSCS III TT&C downlink structure.

A.3 DSCS II TT&C Downlink Section

Figure A.3 shows a simplified block diagram of the DSCS II TT&C downlink section.




Figure A.3 Downlink Section of the DSCS II TT&C Subsystem (19:1)

The discussion will begin with the PCM Mux found on the right-hand side of figure A.3. The Pulse Code Modulated (PCM) Multiplexer is located on the despun section of the DSCS II satellite. It collects and formats all of the telemetry originating on the despun platform. The PCM encoder sends a clock signal to the PCM Mux when it is ready for the PCM Mux to transmit its formatted telemetry to the encoder. There is continuous cross-strapping between the multiplexer and the encoder, which allows either encoder to operate with either multiplexer. The encoder receives the multiplexer data and formats it along with all the other satellite telemetry. Once formatted, the encoder sends the telemetry to both encryptors. The encryptors may be bypassed by ground command, but otherwise all telemetry is encrypted prior to transmission. The selected encrypter will

encrypt the telemetry and pass the encrypted data to the Dual Baseband Assembly Unit (DBAU). The DBAU generates the 1.024 MHz subcarrier and modulates the PCM telemetry onto the subcarrier along with the PRN ranging data passed to the DBAU by the receiver. The signal is then sent from the DBAU to the selected transmitter. The transmitter generates the carrier signal and modulates the subcarrier onto the carrier. The transmitter then transmits the telemetry signal down to the ground antenna. (19:1,2)

A.4 DSCS III TT&C Downlink Section

Figure A.4 shows a simplified block diagram of the DSCS III TT&C downlink section.




Figure A.4 Downlink Section of the DSCS III TT&C Subsystem (21:3)

The large box on the right-hand side of figure A.4 that houses the A and B power supply along with the A and B telemetry units is the Master Telemetry Unit (MTU). The MTU is internally redundant because of the dual power supplies and dual telemetry units. The MTU assembles all of the satellites telemetry using the Analog Multiplexer, Bi-level & Serial/Digital Multiplexer, and Timing & Control. There is complete cross-strapping capabilities between any unit in the A string and any unit in the B string. The MTU also receives telemetry from the Remote Telemetry Unit (RTU). The RTU located on the north panel, collects and processes north panel telemetry. The MTU sends a clock signal to the RTU signaling it to send its telemetry data to the MTU for further processing. Either RTU can operate with either MTU string. Once the telemetry has been formatted it is sent to the selected KG-46 encrypter. There is a encrypter bypass capability that can be accessed by ground command, but has never been used. The encrypter encrypts the telemetry data and passes it to the Baseband Modulator Unit (BMU) housed with the selected transmitter. The BMU generates the 1.024 MHz subcarrier and modulates the received telemetry, along with the PRN data received from the receiver, onto the subcarrier and then passes the signal to the transmitter. The transmitter generates the carrier and modulates the subcarrier onto the carrier and transmits the signal to the ground antenna. (21:3 - 5)

A.5 GPS TT&C Subsystem

The Navstar Global Positioning System (GPS) is a network of navigation satellites placed in a semi-synchronous, 20,200 km circular orbit. The full constellation consists of 24 satellites in six different orbital planes, inclined at 55 degrees. The three-axis stabilized GPS satellite provides position, velocity, and time information to its users worldwide.

GPS has four dedicated ground antennas (GA) located at Diego Garcia, Kwajalein Atoll, Cape Canaveral, and Ascension Island. These ground antennas are used solely by

the 2nd Space Operations Squadron (2SOPS) at Falcon AFB in Colorado to command and control the GPS satellites. The 1st Space Operations Squadron (1SOPS) also performs command and control functions for the GPS satellite bus using the ground antennas of the Air Force Satellite Control Network (AFSCN). The difference between the dedicated antennas and the AFSCN is that the 1SOPS cannot access the satellite payload using the AFSCN ground antenna as the 2SOPS can using the dedicated antennas. The 2SOPS, therefore, has full responsibility for the GPS mission payload. This thesis is not concerned with the operations of the mission payload because it is specific to GPS satellites and could not be controlled generically. The GPS TT&C subsystem will now be described and those areas specific to the use of the dedicated ground antennas will be ignored in this discussion.

Figure A.5 shows a simplified block diagram of the GPS Block II TT&C subsystem. It includes both the uplink and downlink sections of the GPS TT&C subsystem. The uplink section will be described first, followed by the function of the downlink section.

A.5.1 GPS TT&C Uplink Section. Beginning in the top left corner of the diagram, the GPS satellite has three TT&C antennas. There are two conical and one biconical antenna. The two conical antennas are located on separated faces of the satellite. This configuration of three antennas ensures the satellite control center will be able to establish a communication link with a satellite in any orientation. The biconical and aft conical antennas are used mostly during launch and early orbit, while the forward conical antenna is used during nominal on-orbit operations. The diplexer within the Radio Frequency (RF) Assembly allows simultaneous receive and transmit capability and the hybrid splits the power of the uplinked signal in half and routes half to each receiver. (7:4)




Figure A.5 The GPS TT&C Subsystem

The receiver/demodulators are continuously on and are always searching for an uplink carrier with a frequency of 1783.74 MHz (SGLS Ch. 6). When this frequency is present, both receivers will lock onto the uplinked signal. There is a component located within the receiver/demodulator that detects modulation on the uplink. When modulation is present on the uplink this component will recognize the modulation and send a Decoder/Activate (D/A) signal to the front end of the Dual Command Decoder (DCD). The D/A signal notifies the DCD that a modulated uplink signal has been detected. At the presence of a D/A signal, the DCD powers both of the KIR decrypters in anticipation of a command. Both receivers will demodulate the uplinked command and strip any PRN ranging data off of the uplinked carrier, if present. The command is converted into 1 and 0 tones. These 1 and 0 command tones, along with uplinked S-pulses and a clock signal are sent from both receivers to the front end of their respective DCD. The DCD conditions the signals and passes them to the corresponding KIRs. The KIR which matches the encryption scheme of the uplinked command will authenticate the signal by comparing an uplinked Ground Command Count (GCC) against the internal Vehicle Command Count (VCC). If the two seven digit numbers match, the command is authenticated and decrypted by the KIR decrypter. Once decrypted the KIR sends the command to the back-end of the DCD for processing. (7:5)

The DCD is internally redundant. It has both an A and a B side and, as stated in the above paragraphs, a front and back side. The back side is dormant until a command is received from a KIR decrypter. Once received, the back side of the DCD becomes active and performs several internal checks on the command before executing the command. If the A side of the DCD is active and the satellite operator sends a command into the B side of the DCD, the command will be ignored. If the operator wishes to use the secondary command path, he or she must first send a command to turn on the B side of the DCD. (2)

In a nominal configuration, receiver one output is sent through the primary KIR decrypter and receiver two output is sent through the secondary KIR. The only way to route the output of receiver one to the secondary KIR and receiver two output to the primary KIR is by automatic or commanded cross-strapping. Within the DCD is a cross-strap timer. If a command is not processed by the satellite within six days, the cross-strap timer will activate and will automatically cross-strap the receivers and KIR components. This cross-strap capability can also be accessed by ground command. The next component shown in the above diagram is the Navigation Data Unit (NDU). The NDU is part of the navigation mission package and will not be discussed. That completes the discussion of the GPS TT&C uplink section. Next is the discussion of the functions of the GPS TT&C downlink section. (7:10)

A.5.2 GPS TT&C Downlink Section. The Pulse Code Modulator (PCM) encoder collects all of the satellite telemetry, processes and combines it into one continuous telemetry stream. The PCM is internally redundant with both an A and B side. Each side of the PCM consists of seven components. Five of these components are completely cross-strapable with components on the other side of the PCM. This cross-strap capability allows for 192 possible PCM configurations. Each side of the PCM is hard-wired to its corresponding encrypter. There is no cross-strap capability between the PCM and the encrypters. The telemetry stream created within the PCM encoder is sent into the KG encrypter. The KG encrypts the telemetry and sends the data to both transmitters. The GPS TT&C subsystem has no capability to bypass its encrypters. (2)

The powered transmitter will generate a 1.7 MHz subcarrier and modulate the telemetry stream onto the subcarrier. It will then generate a carrier and modulate the subcarrier onto the carrier. The carrier is passed to the RF assembly for transmission through the selected antenna to the ground. (2)

APPENDIX B. Data Representation Schemes

B.1 GISMO Variable/Qualifier Format

B.1.1 Actual Telemetry. All variable names used to hold actual satellite state values begin with an "A". For example: AINYS is the variable name for the actual telemetry reading of the INYS sensor point.

B.1.2 Expected Telemetry. All variable names used to hold expected satellite state values are given the same name as the satellite state variable. For example: INYS is the variable name for the expected value of the satellite state variable, INYS. For a floating point, there is no expected telemetry value so this expected variable would hold a Failed/Working status of the component represented by the variable. For example: RC1P15 is the variable name for receiver converter one and will hold either the value Failed or Working depending on the status of the receiver converter.

B.1.3 Availability of Redundant Units. Along with the expected status of each satellite state variable, there is a variable used to represent the redundancy of the component represented by the state variable. This is not used for floating point values. The format adds an RD extension to the original name of the state variable. For example: INYS_RD is the variable name used within GISMO to hold the redundancy status (YES or NO) of the INYS component.

B.1.4 Qualifiers to Represent Out-of-Limit Conditions. Both for the floating point telemetry values and the discrete telemetry values, a qualifier is used to

represent their OOL condition. The qualifier within EXSYS is given the same name as the state variable and the possible values the qualifier could be set to are as follows:

- 1 below the lower red limit
- 2 below the lower yellow limit
- 3 nominal
- 4 above the upper yellow limit
- 5 above the upper red limit
- 6 abnormal

A qualifier for a floating point would be set to any of the first 5 possibilities shown above, while the discrete state variables would be represented using either the nominal or abnormal settings.

APPENDIX C. Detailed Analysis

This appendix presents the details of the telemetry comparison made between the DSCS II and DSCS III satellites, and the comparison made between the DSCS II, DSCS III and GPS satellites. After the telemetry comparison analysis is shown, a detailed analysis of the detectable TT&C anomalies for all three satellites are presented.

C.1 DSCS II and DSCS III Telemetry Comparison

DSCS II TT&C Telemetry	DSCS III TT&C Telemetry	GENERIC	Brief Component Description
IRON	VEHID	IRON	satellite operations number
SIGPR	RCVRLK	RCVRLK	receiver lock status
SIGCON	AMSYNC	TONES	tone detector status
INYS	INY	INYS	decrypter status
AUTH1S	AUTH1S	AUTH1S	primary INY authentication
AUTH2S	AUTH2S	AUTH2S	secondary INY authentication
VCC1	VCC1	VCC1	vehicle command count from INY1
VCC2	VCC2	VCC2	vehicle command count from INY2
CMD1CK	CMDVER		
CMD2CK			
CMDLFF			
	CMDW		
	CMDDEC		
	MSGMD		
	CMDCNT		
EIACBS			
EIABSI			
EIAT			
EIA1+V			
EIA2+V			
EIAVOL			
EIAMAT			
TTCBSI	TTCI	TTCBSI	TT&C bus current
MNBUSV	MNBV	MNBUSV	main bus voltage
MNBUSI			
RCVRT			
RC1+15			
RC2+15			
RCV1SF	RCVALS	RCVASF/LS	primary receiver sweep freq/loop stress
	RCVASS		
RCV2SF	RCVBLS	RCVBBSF/LS	secondary receiver sweep freq/loop stress
	RCVBSS		
	RCVINY		
	K23A28		
	K23B28		
ENCRP	K46A28		
	K46B28		
	KG46BY		
ENCRPT			
	AUTODL		

DSCS II TT&C Telemetry	DSCS III TT&C Telemetry	GENERIC	Brief Component Description
	SBDL		
TMX	TMX	TMX	transmitter status
TXCOHO	TXCOHO	TXCOHO	coherent mode status
TMXTP	TMXAP		
	TMXBP		
TMX_T	TMXAT		
	TMXBT		
{automatic KG/XMT xstrp}	TXATLM		
{automatic KG/XMT xstrp}	TXBTLM		
	TXARNG		
	TXBRNG		
TMX+15			
MPX	RTU	RMPX	remote multiplexer
	RTUASV		
	RTUBSV		
MPXT	RTUT	RMPXT	remote multiplexer temperature
	NPCONV		
	NPCOFF		
	NPC+24		
	NPC+15		
	NPC-15		
	NPC+12		
	NPC-12		
	NPC+8		
	NPC+5		
	TLMPWR		
	TLMOFF		
EIAT			
MXCALV			
	SPCONV		
	SPCOFF		
	SPC+28		
	SPC+5		
	SPCSV1		
	SPCSV2		
	SPCSV3		
	MTUSV		
	CDASV		
	CDBSV		

DSCS II TT&C Telemetry	DSCS III TT&C Telemetry	GENERIC	Brief Component Description
-----	ANLMUX	-----	
-----	DIGMUX	-----	
ENCDR	CTULOG	ENCDR	encoder status
ENCDRT	CTUT	ENCDRT	encoder temperature
ENCALV	-----	-----	
SLA	-----	-----	
SLA1+V	-----	-----	
SLA2+V	-----	-----	
SLAVOL	-----	-----	
SLAMAT	-----	-----	
SLABSI	-----	-----	
Telemetry Generated on the Ground:			
SPE	SPE	SPE	S-pulse status
GCC1	GCC1	GCC1	ground command count for INY1
GCC2	GCC2	GCC2	ground command count for INY2
ACTIVE/PASSIVE	ACTIVE/PASSIVE	ACTV_/PASV	ground antenna status
AZ	AZ	AZ	ground antenna azimuth angle
EL	EL	EL	ground antenna elevation angle
RANGE	RANGE	RANGE	range of satellite from earth

C.2 GPS Telemetry Comparison

LEGEND: ! - This telemetry point was common to DSCS II and DSCS III, but is not common to GPS.

DSCS II Telemetry	DSCS III Telemetry	GPS Telemetry	GENERIC Telemetry
IRON	VEHID	IRON	IRON
SIGPR	RCVRLK	CARPRE	RCVRLK
SIGCON	AMSYNC	DARCVR	TONES
INYS	INY	CMSC1	INYS
		CMSC2	
AUTH1S	AUTH1S	CMSC1	AUTH1S
AUTH2S	AUTH2S	CMSC2	AUTH2S
VCC1	VCC1	CMSC1	VCC1
VCC2	VCC2	CMSC2	VCC2
CMD1CK	CMDVER	DCDAST	CMDCK
CMD2CK		DCDBST	-----
CMDLFF	-----	-----	-----
-----	CMDW	-----	-----
-----	CMDDEC	DCDSEL	-----
-----	-----	DISOD	-----
-----	-----	LOGA5V	-----
-----	-----	LOGB5V	-----
-----	-----	DISA5V	-----
-----	-----	DISB5V	-----
-----	MSGMD	-----	-----
-----	CMDCNT	-----	-----
EIACBS	-----	-----	-----
EIABSI	-----	-----	-----
EIAT	-----	-----	-----
EIA1+V	-----	-----	-----
EIA2+V	-----	-----	-----
EIAVOL	-----	-----	-----
EIAMAT	-----	-----	-----
TTCBSI	TTCI	-----	!!! TTCBSI !!!
MNBUSV	MNBV	-----	!!! MNBUSV !!!
MNBUSI	-----	-----	-----
RCVRT	-----	RCVR1T	-----
-----	-----	RCVR2T	-----
RC1+15	-----	-----	-----
RC2+15	-----	-----	-----
RCV1SF	RCVALS	-----	!!! RCVASF/LS !!!
-----	RCVASS	-----	-----
RCV2SF	RCVBLS	-----	!!! RCVBSF/LS !!!
-----	RCVBSS	RCVR1L	-----
-----	-----	RCVR2L	-----

DSCS II Telemetry	DSCS III Telemetry	GPS Telemetry	GENERIC Telemetry
-----	RCVINY	XSTRP	-----
-----	-----	CLRENA	-----
-----	K23A28	-----	-----
-----	K23B28	-----	-----
ENCRP	K46A28	KG1PSV	-----
-----	K46B28	KG2PSV	-----
-----	KG46BY	-----	-----
ENCRPT	-----	-----	-----
-----	AUTODL	ATOCON	-----
-----	SBDL	-----	-----
TMX	TMX	XMTSEL	TMX
TXCOHO	TXCOHO	-----	!!! *TXCOHO !!!
TMXTP	TMXAP	XMT1P	-----
-----	TMXBP	XMT2P	-----
TMX_T	TMXAT	XMT1T	-----
-----	TMXBT	XMT2T	-----
-----	TXATLM	-----	-----
-----	TXBTLM	-----	-----
-----	TXARNG	-----	-----
-----	TXBRNG	-----	-----
TMX+15	-----	-----	-----
MPX	RTU	-----	!!! RMPX !!!
-----	RTUASV	-----	-----
-----	RTUBSV	-----	-----
MPXT	RTUT	-----	!!! RMPXT !!!
-----	NPCONV	-----	-----
-----	NPCOFF	-----	-----
-----	NPC+24	-----	-----
-----	NPC+15	-----	-----
-----	NPC-15	-----	-----
-----	NPC+12	-----	-----
-----	NPC-12	-----	-----
-----	NPC+8	-----	-----
-----	NPC+5	-----	-----
-----	TLMPWR	-----	-----
-----	TLMOFF	-----	-----
EIAT	-----	-----	-----
MXCALV	-----	-----	-----

DSCS II Telemetry	DSCS III Telemetry	GPS Telemetry	GENERIC Telemetry
-----	SPCONV	-----	-----
-----	SPCOFF	-----	-----
-----	SPC+28	-----	-----
-----	SPC+5	-----	-----
-----	SPCSV1	-----	-----
-----	SPCSV2	-----	-----
-----	SPCSV3	-----	-----
-----	MTUSV	-----	-----
-----	CDASV	-----	-----
-----	CDBSV	-----	-----
-----	ANLMUX	A/DCON	-----
-----	DIGMUX	DIGITL	-----
-----	-----	CONLOG	-----
-----	-----	DIFFA	-----
-----	-----	DRATE	-----
-----	-----	FORMAT	-----
-----	-----	N/GD	-----
-----	-----	COMB	-----
ENCDR	CTULOG	PCMPE	ENCDR
-----	-----	PCMPWR	-----
ENCDRT	CTUT	-----	!!! ENCDRT !!!
ENCALV	-----	PCMC3	-----
-----	-----	PCMC38	-----
-----	-----	PCM+12	-----
-----	-----	PCM-12	-----
-----	-----	PCM+5	-----
SLA	-----	-----	-----
SLA1+V	-----	-----	-----
SLA2+V	-----	-----	-----
SLAVOL	-----	-----	-----
SLAMAT	-----	-----	-----
SLABSI	-----	-----	-----
-----	-----	ANTSEL	-----
-----	-----	ANTDEP	-----
-----	-----	XSTIM	-----
-----	-----	TIMOPR	-----
-----	-----	ATOPS	-----
-----	-----	ATOPCM	-----
-----	-----	ATOXMT	-----
-----	-----	XMTATO	-----
-----	-----	NCONF	-----
-----	-----	NBUFF	-----
-----	-----	CLKENA	-----
Telemetry Generated on the Ground:			
SPE	SPE	SPE	SPE
GCC1	GCC1	G PRI	GCC1
GCC2	GCC2	G SEC	GCC2
ACTIVE/PASSIVE	ACTIVE/PASSIVE	-----	ACTV_PASV
AZ	AZ	AZ	AZ
EL	EL	EL	EL
RANGE	RANGE	RANGE	RANGE

C.3 Detectable and Undetectable TT&C Anomalies

C.3.1 DSCS II Anomalies.

<u>Anomaly</u>	<u>Detectable by GISMO?</u>
GENERIC ANOMALIES:	
Failed SIGCON	Yes. Assumes, S-pulses have been cycled
Failed Receiver	Yes. Assumes uplink modulation is cycled
Failed INY	Yes.
Invalid Range Contingency	Yes.
TT&C#1 - No Carrier Present (failed transmitter)	Yes.
TT&C#2 - Carrier, no subcarrier (failed subcarrier generator within transmitter)	Yes.
TT&C#3 - Carrier, subcarrier, no tlm mod. (failed encoder or encrypter or transmitter)	Yes. One level diagnostic depth: There are three major components that could fix this anomaly, but GISMO has no memory capability so it swaps the most likely unit to be causing the problem and assumes the problem is fixed. In a real world scenario, the expert system will need to remember what is swapped and if that does not fix the problem take another action to fix the problem.
TT&C#4 -Car., subcar., mod., no synch (failed encoder or encrypter or transmitter)	Yes. One level diagnostic depth: There are three major components that could fix this problem, but GISMO has no memory capability so it swaps the most likely unit to be causing the problem and assumes the problem is fixed. In a real world scenario, the expert system will need to remember what is swapped and if that does not fix the problem take another action to fix the problem.
Multiplexer failure	Yes, using satellite specific rules This anomaly is very similar to the DSCS III Remote Telemetry Unit (RTU) failure. The DSCS III RTU and the DSCS II multiplexer perform similar functions and were selected as common components during the one-to-one telemetry comparison analysis. The differences lie in the method of detecting a failure of these common components. These differences prevented detection using generic rules.

<u>Anomaly</u>	<u>Detectable by GISMO?</u>
----------------	-----------------------------

DSCS II SPECIFIC ANOMALIES:

Failed Receiver Converter	Yes, using satellite specific rules
Command Check Glitch	Yes, using satellite specific rules
SLA Matrix Glitch	Yes, using satellite specific rules
SLA Path Glitch	Yes, using satellite specific rules
EIA Matrix Glitch	Yes, using satellite specific rules

C.3.2 DSCS III Anomalies.

<u>Anomaly</u>	<u>Detectable by GISMO?</u>
----------------	-----------------------------

GENERIC ANOMALIES:

Failed AMSYNC	Yes. Assumes, S-pulses have been cycled
Failed Receiver	Yes. Assumes uplink modulation is cycled
Failed INY	Yes.
Invalid Range Contingency	Yes.
TT&C#1 - No Carrier Present (failed transmitter)	Yes.
TT&C#2 - Carrier, no subcarrier (failed subcarrier generator in transmitter)	Yes.
TT&C#3 - Carrier, subcarrier, no tlm mod. (failed encoder or encrypter or transmitter)	Yes. One level diagnostic depth: There are three major components that could fix this anomaly, but GISMO has no memory capability so it swaps the most likely unit to be causing the problem and assumes the problem is fixed. In a real world scenario, the expert system will need to remember what is swapped and if that does not fix the problem take another action to fix the problem.
TT&C#4 - Car., subcar., mod, no synch (failed encoder or encrypter or transmitter)	Yes. One level diagnostic depth: There are three major components that could fix this anomaly, but GISMO has no memory capability so it swaps the most likely unit to be causing the problem and assumes the problem is fixed. In a real world scenario, the expert system will need to remember what is swapped and if that does not fix the problem take another action to fix the problem.
Abnormal Telemetry (failed RTU)	No. This anomaly is very similar to the DSCS II multiplexer anomaly. The DSCS III Remote Telemetry Unit (RTU) and the DSCS II multiplexer perform a common functions and were selected as common components during the one-to-one telemetry comparison analysis. The differences lie in the method of detecting a failure of these common components. These differences prevented detection using generic rules.

Anomaly

Detectable by GISMO?

DSCS III SPECIFIC ANOMALIES:

ACE Command Port Clear	No. Anomaly occurs while commanding
ACE Command Timing Rejection Clear	No. Anomaly occurs while commanding
Command Verification But No Functional Verification:	No. Anomaly occurs while commanding
CD not Accepting Commands	No. Anomaly occurs while commanding
Command Path Anomaly	No. Anomaly occurs while commanding

C.3.3 GPS Anomalies.

<u>Anomaly</u>	<u>Detectable by GISMO?</u>
----------------	-----------------------------

GENERIC ANOMALIES:

Failed DARCVR (Assumes S-pulses have been cycled)	COM mode - Yes. ATO mode - No. CTO mode - No.
Failed CARPRE (Assumes uplink modulation is cycled)	COM mode - Yes. ATO mode - No. CTO mode - No.
Anomalous VCC (failed INY)	COM mode - No. ATO mode - No. CTO mode - No.
TT&C#1 - No Carrier Present (failed transmitter)	COM mode - Yes. ATO mode - Yes. CTO mode - No.
TT&C#2 - Carrier, no subcarrier (failed subcarrier generator in transmitter)	COM mode - Yes. ATO mode - Yes. CTO mode - Yes.
TT&C#3 - Carrier, subcarrier, no tlm mod. (failed encoder or encrypter or transmitter)	COM mode - Yes. ATO mode - Yes. CTO mode - Yes.

One level diagnostic depth: There are three major components that could fix this anomaly, but GISMO has no memory capability so it swaps the most likely unit to be causing the problem and assumes the problem is fixed. In a real world scenario, the expert system will need to remember what is swapped and if that does not fix the problem take another action to fix the problem.

TT&C#4 - Car., subcar., mod, no synch (failed encoder or encrypter or transmitter)	COM mode - Yes. ATO mode - Yes. CTO mode - Yes.
---	---

One level diagnostic depth: There are three major components that could fix this anomaly, but GISMO has no memory capability so it swaps the most likely unit to be causing the problem and assumes the problem is fixed. In a real world scenario, the expert system will need to remember what is swapped and if that does not fix the problem take another action to fix the problem.

Anomaly

Detectable by GISMO?

GPS SPECIFIC ANOMALIES:

Failure to switch bicone to conical antennas	No.
Failure to switch conical to bicone antenna	No.
Primary receiver failure - ATO verification	No.
Unable to reconfigure primary PCM	No.
Over-Modulated turned-around command tones	No.
Commanding Contingency Flow Diagram (Anomaly occurs while commanding)	No.

Bibliography

1. Bleier, S., S. Hollander, and S. Sutton. "Advanced Satellite Workstation (ASW)." *AGARD Conference Proceedings 499: Machine Intelligence for Aerospace Electronic Systems*. May 1991.
2. Brewer, Jim. GPS TT&C specialist, Rockwell International, Colorado Springs CO. Telephone interview. 28 September 1994.
3. Busse, Carl. "An Expert System for Satellite and Instrument Data Anomaly and Fault Isolation." *NASA, Marshall Space Flight Center, Fourth Conference on Artificial Intelligence for Space Applications*. October 1988.
4. Contreras, Vincent J. and Jay C. Ferguson. "Knowledge Management." *MILCOM IEEE Military Communications Conference*. October 1986.
5. Drabble, Brian. "Spacecraft Command and Control using Artificial Intelligence Techniques." *Journal of the British Interplanetary Society New Space Concepts (Part III)*. Vol.44 No. 6 June 1991.
6. Ferguson, Jay C. "STAR-PLAN: A Satellite Anomaly Resolution and Planning System." *Coupling Symbolic and Numerical Computing in Expert Systems*. 1986.
7. First Satellite Operations Squadron. "Telemetry, Tracking, and Commanding." *GPS Orbital Operations Handbook*, Volume I, Section 2.2. Revision A.
8. Golden, Marilyn, Ronald W. Siemens, and Jay C. Ferguson. "What's Wrong With Rules?" Ford Aerospace & Communications Corporation Internal Report. Westex 1986.
9. Golden, Marilyn. Manager of Modeling and Simulation, Systems Engineering Laboratory, Loral Space and Range Systems, Sunnyvale Ca. Telephone interview. 11 October 1994.

10. Hartman, Wayne. "Challenges in Building Intelligent Systems for Space Mission Operations." *Fourth Annual Workshop on Space Operations Applications and Research (SOAR 90)*, Volume I, Albuquerque, New Mexico. June 1990.
11. Heher, Dennis and Paul Pownall. "StarPlan: A Model-based Diagnostic System for Spacecraft." *Space Operations Automation & Robotics Workshop*. July 1989.
12. Integrated Satellite Control Standards Management Committee briefing. 3-4 December 1992.
13. Quinn, Todd and Richard Schlegelmilch. "Applying Model-Based Diagnostics to Space Power Distribution." *Knowledge-Based Artificial Intelligence Systems in Aerospace and Industry*. April 1994.
14. Schwuttke, U.M., A.G. Quan, and B. Holder. "Intelligent Data Presentation for Real-time Spacecraft Monitoring." *Applications of Artificial Intelligence 1993: Knowledge-Based Systems in Aerospace and Industry*. April 1993.
15. Skinner, James M., and George F. Luger. "A Synergistic Approach to Reasoning for Autonomous Satellites." *Proceedings of the Advisory Group for Aerospace Research and Development (AGARD) Panel on Machine Intelligence for Aerospace Electronic Systems*. May 1991.
16. Skinner, James M. and George F. Luger. "An Integrated Reasoner for Diagnosis in Satellite Control." *Applications of Artificial Intelligence 1993: Knowledge-Based Systems in Aerospace and Industry*. April 1993.
17. Turner, S.R., P.K. Mangan, M.D. Melon, and C.B. Simmons. "Prototype Expert System Workstation for Satellite Anomaly Resolution." *Journal of Spacecraft and Rockets*. May-June 1991.
18. Turner, Scott R. and Patricia Mangan. "Reusable Expert Systems." *The Fourth International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems, Kauai, Hawaii*. 1991.

19. Third Satellite Operations Squadron. "Satellite TT&C Subsystem." *DSCS II Operations Manual*, Volume I, Section 5-2. September 1992.
20. Third Satellite Operations Squadron. "Mission of DSCS II." *DSCS II Operations Manual*, Volume I, Section 4-1. September 1992.
21. Third Satellite Operations Squadron. "Telemetry, Tracking & Commanding." *DSCS III Operations Manual*, Volume I, Section 5-2. January 1993.
22. Wagner, R.E. and A.N. Blasdel. "An Approach to Autonomous Attitude Control for Spacecraft." *Guidance and Control 1988, Volume 66: Advances in the Astronautical Sciences*. American Astronautical Society 1988.
23. Wilson, Andrew. *Jane's Space Directory 1993-94* (ninth edition). Great Britain: Jane's Data Division, 1993.

Vita

Captain Loretta A. Kelemen was born on 25 January 1966 in Marietta, Georgia. She graduated from Clarkston High School in Clarkston, Georgia in 1984. She received her undergraduate degree in Computer Science from Valdosta State College in Valdosta, Georgia in June 1988. In addition to the award of distinguished graduate, Captain Kelemen received a regular commission in the USAF. Her first assignment was to the 3rd Satellite Operations Squadron (3SOPS), Falcon AFB, Colorado, as a Defense Satellite Communication System Phase II (DSCS II) planner analyst. In 1991, she became the chief DSCS II instructor, responsible for qualification and recurring training of all DSCS II satellite operations crew commanders. In 1992, Captain Kelemen was selected to be the chief DSCS II evaluator, responsible for evaluating the skill and knowledge of all DSCS II satellite operations crew commanders. She was selected to attend AFIT in 1993 and following graduation, will be assigned to the 45th Range Squadron at Cape Canaveral, Florida.

Permanent Address: 1075 Singley Drive
Locust Grove, GA 30248

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE December 1994	3. REPORT TYPE AND DATES COVERED Master's Thesis		
4. TITLE AND SUBTITLE GENERIC SATELLITE MONITORING EXPERT SYSTEM		5. FUNDING NUMBERS		
6. AUTHOR(S) Loretta A. Kelemen, Captain, USAF				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology, WPAFB OH 45433-6583		8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GSO/ENG/94D-02		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Nancy L. Crowley, Lt Col, Deputy Chief Satellite Control and Simulation Division PL/VIQ, 3550 Aberdeen Avenue SE Kirtland AFB, NM 87117-5776		10. SPONSORING / MONITORING AGENCY REPORT NUMBER		
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Distribution Unlimited		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) Air Force Satellite Operations is undergoing major changes. Operators no longer receive detailed satellite training, instead they are taught basic fundamentals of satellite operations and expected to control multiple multi-million dollar satellites. The need is clear. An efficient and economical automated system is necessary to assist the satellite operator in his daily tasks of maintaining these DOD priority resources. Satellite intelligent controllers have been under R&D since the early 1980s to meet this need. These systems, however, have focused on the control of one constellation of satellites. In a military striving for efficiency and lower costs, developing a unique intelligent controller for each satellite constellation is unaffordable. This research provided support for the concept of a generic satellite intelligent controller, through the development of a prototype expert system. This capability would allow a generic rule-base to operate and maintain multiple satellite systems. The initial prototype detected anomalies on one subsystem of two different satellites. After the two-satellite prototype was created, a third satellite was analyzed to show support for the viability of the two-satellite prototype. More research is necessary, but this thesis has created support for the concept of generic satellite controllers and has laid the foundation for future extensions.				
14. SUBJECT TERMS ARTIFICIAL INTELLIGENCE, SATELLITE GROUND SUPPORT COMMAND AND CONTROL, EXPERT SYSTEMS, SATELLITE OPERATIONS, SATELLITE INTELLIGENT CONTROLLERS, KNOWLEDGE BASED SYSTEMS		15. NUMBER OF PAGES 129		
		16. PRICE CODE		
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

GENERAL INSTRUCTIONS FOR COMPLETING SF 298

The Report Documentation Page (RDP) is used in announcing and cataloging reports. It is important that this information be consistent with the rest of the report, particularly the cover and title page. Instructions for filling in each block of the form follow. It is important to *stay within the lines* to meet *optical scanning requirements*.

Block 1. Agency Use Only (Leave blank).

Block 2. Report Date. Full publication date including day, month, and year, if available (e.g. 1 Jan 88). Must cite at least the year.

Block 3. Type of Report and Dates Covered. State whether report is interim, final, etc. If applicable, enter inclusive report dates (e.g. 10 Jun 87 - 30 Jun 88).

Block 4. Title and Subtitle. A title is taken from the part of the report that provides the most meaningful and complete information. When a report is prepared in more than one volume, repeat the primary title, add volume number, and include subtitle for the specific volume. On classified documents enter the title classification in parentheses.

Block 5. Funding Numbers. To include contract and grant numbers; may include program element number(s), project number(s), task number(s), and work unit number(s). Use the following labels:

C - Contract	PR - Project
G - Grant	TA - Task
PE - Program Element	WU - Work Unit Accession No.

Block 6. Author(s). Name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. If editor or compiler, this should follow the name(s).

Block 7. Performing Organization Name(s) and Address(es). Self-explanatory.

Block 8. Performing Organization Report Number. Enter the unique alphanumeric report number(s) assigned by the organization performing the report.

Block 9. Sponsoring/Monitoring Agency Name(s) and Address(es). Self-explanatory.

Block 10. Sponsoring/Monitoring Agency Report Number. (If known)

Block 11. Supplementary Notes. Enter information not included elsewhere such as: Prepared in cooperation with...; Trans. of...; To be published in.... When a report is revised, include a statement whether the new report supersedes or supplements the older report.

Block 12a. Distribution/Availability Statement. Denotes public availability or limitations. Cite any availability to the public. Enter additional limitations or special markings in all capitals (e.g. NOFORN, REL, ITAR).

DOD - See DoDD 5230.24, "Distribution Statements on Technical Documents."

DOE - See authorities.

NASA - See Handbook NHB 2200.2.

NTIS - Leave blank.

Block 12b. Distribution Code.

DOD - Leave blank.

DOE - Enter DOE distribution categories from the Standard Distribution for Unclassified Scientific and Technical Reports.

NASA - Leave blank.

NTIS - Leave blank.

Block 13. Abstract. Include a brief (*Maximum 200 words*) factual summary of the most significant information contained in the report.

Block 14. Subject Terms. Keywords or phrases identifying major subjects in the report.

Block 15. Number of Pages. Enter the total number of pages.

Block 16. Price Code. Enter appropriate price code (*NTIS only*).

Blocks 17. - 19. Security Classifications. Self-explanatory. Enter U.S. Security Classification in accordance with U.S. Security Regulations (i.e., UNCLASSIFIED). If form contains classified information, stamp classification on the top and bottom of the page.

Block 20. Limitation of Abstract. This block must be completed to assign a limitation to the abstract. Enter either UL (unlimited) or SAR (same as report). An entry in this block is necessary if the abstract is to be limited. If blank, the abstract is assumed to be unlimited.