

**SHARP: A MULTI-MISSION AI SYSTEM
FOR SPACECRAFT TELEMETRY MONITORING AND DIAGNOSIS**

Denise L. Lawson
Mark L. James

*Artificial Intelligence Group
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109*

ABSTRACT

The Spacecraft Health Automated Reasoning Prototype (SHARP) is a system designed to demonstrate automated health and status analysis for multi-mission spacecraft and ground data systems operations. Telecommunications link analysis of the Voyager II spacecraft is the initial focus for the SHARP system demonstration which will occur during Voyager's encounter with the planet Neptune in August, 1989, in parallel with real-time Voyager operations.

The SHARP system combines conventional computer science methodologies with artificial intelligence techniques to produce an effective method for detecting and analyzing potential spacecraft and ground systems problems. The system performs real-time analysis of spacecraft and other related telemetry, and is also capable of examining data in historical context.

This paper gives a brief introduction to the spacecraft and ground systems monitoring process at the Jet Propulsion Laboratory (JPL). It describes the current method of operation for monitoring the Voyager Telecommunications subsystem, and highlights the difficulties associated with the existing technology. The paper

details the approach taken in the SHARP system to overcome the current limitations, and describes both the conventional and artificial intelligence solutions developed in SHARP.

INTRODUCTION

The Voyager II spacecraft was launched from Cape Canaveral, Florida on August 20, 1977. The technology to track and monitor such a probe was designed and developed in the early 1970's. This now antiquated technology, coupled with the efforts of bright, resourceful scientists, has carried Voyager through near-fatal catastrophic events to three of our solar system's outer planets (four in August). Despite failed radio receivers, sunlight damage to the photopolarimeter scientific instrument, and a partially paralyzed scan platform (which houses Voyager's imaging system), these scientists have kept Voyager operational, enabling the capture and transmission of vast amounts of invaluable information and images of the Jovian, Saturnian, and Uranian systems.

During critical periods of the mission, up to 30 real-time operators are required to monitor the spacecraft's ten subsystems on a 24-hour, 7-day per week schedule. This does not include the numerous subsystem and scientific

instrument specialists that must constantly be available on call to handle emergencies.

As more and more solar system explorations are undertaken, it will become increasingly difficult to staff a large enough effort to support these expensive missions. Currently there is one Mission Control Team and one Spacecraft Team for each flight project. JPL has initiated an effort to coordinate all missions through a central Space Flight Operations Center (SFOC) whose goal is to transition from single-project dedicated flight teams to one multi-mission team which flies all spacecraft. Within SFOC, the Voyager spacecraft will continue to be monitored throughout their extended mission of discovering the solar system heliopause (the boundary between the Sun's magnetic influence and interstellar space); the Magellan spacecraft, to be launched in April, 1989, will be tracked throughout its flight to Venus; the Galileo mission to Jupiter will be monitored; and other new flight projects will be observed throughout their operation by this single multi-mission flight team.

The Spacecraft Health Automated Reasoning Prototype (SHARP) is an effort to apply artificial intelligence (AI) techniques to the task of multi-mission monitoring of spacecraft and diagnosis of anomalies. Ultimately, SHARP will ease the burden that multiple missions would inevitably place upon subsystem, scientific instrument, and Deep Space Network (antenna) experts. SHARP will automate many of the mundane analysis tasks, and reduce the number of operators required to perform real-time monitoring activities. The system will enhance the reliability of monitoring operations, and may prevent those types of errors that cause spacecraft such as the Soviet Phobos to be lost.

The Voyager II spacecraft was targeted for the prototype effort since it is currently the only spacecraft in flight

which has yet to complete its primary mission. The prototype effort was further focused to one subsystem so that specific concepts could be developed and then demonstrated in a vigorous operational setting: the spacecraft's encounter of the planet Neptune in August, 1989. The Telecommunications subsystem was chosen for the initial demonstration since anomalies occur on a frequent basis in this area, and the Telecom expert, Mr. Boyd Madsen, demonstrated enthusiastic support. The Telecommunications area also presents the challenge of coordinating monitoring and diagnosis efforts of both the spacecraft and ground data systems.

Like many artificial intelligence applications, in order to supply the AI component of a system with real data, a substantial effort must be invested in the development of other aspects of the system. This may entail utilizing standard conventional computer science methodologies and enhanced graphical capabilities. The SHARP system efficiently incorporates these technologies to complement the use of AI techniques.

CURRENT METHOD

In current mission operations practice each spacecraft is monitored daily, and during planetary encounters monitoring is continuous. Three complexes of antennas located around the world comprise NASA's Deep Space Network (DSN), in the Mojave desert at Goldstone, California; near Canberra, Australia; and near Madrid, Spain. With the exception of occultations and a short gap between the Canberra and Madrid stations, the spacecraft is always in view from one of these Deep Space Stations (DSS). Such a scheduled period of observance of the spacecraft by a DSS is called a pass.

Required Data

In order to effectively analyze the telecommunications link from the

spacecraft through the Deep Space Network and ultimately to the computers at JPL, a wide variety of information must be accessed and processed. This analysis occurs in real-time as well as prior to the scheduled spacecraft pass.

Predicts are numerical predictions of acceptable threshold values for particular spacecraft and DSS parameters. The current method of generating pass predicts is by searching large hardcopy listings of raw predicts to find the correct spacecraft, station, time, and other approximated information. Predicts are then manually corrected to reflect the actual spacecraft state using a hand calculator, and the results are manually recorded on a data sheet.

Another piece of information pertinent to spacecraft monitoring is the Integrated Sequence of Events (ISOE). The ISOE is a hardcopy of scheduled spacecraft activity integrated with DSS information. ISOE data is used extensively throughout the monitoring process in predict data, alarm determination, graphics, and diagnosis. The Voyager ISOE must be visually scanned, and Telecom events manually highlighted by the real-time operator so that the Telecom activity can be monitored. A handwritten correction sheet is issued for each modification to an ISOE.

Telemetry data from the spacecraft, tracking stations, and other relevant systems is collected in the JPL computers and separated into channels that are distributed across JPL for processing and analysis. These channels contain the values of hundreds of spacecraft engineering parameters and station performance parameters. The channels are plotted on black and white computer screens and are visually monitored to ensure that they remain within their pre-specified limits.

Also critical to the communications link analysis are alarm limits, the threshold values for spacecraft and DSS

performance. These values are selected according to the status of several parameters. However, the process to change these limits is manual and must be performed in real-time. The procedure is so impeditive, and occurs so often, that typically a wide threshold is selected which incorporates the entire range of parameter conditions, creating the risk of undetected anomalies.

Limitations

Due to cumbersome and time-consuming processes, several limitations exist on the current method of analyzing Voyager telecommunications link data.

The tedious manual process for predict generation may take up to two hours each day, and limits calculations to one predict point per hour. Actual link parameters may be received every 15 seconds, leaving quite a disparity between the desired number of predictions and the incoming data.

The Integrated Sequence of Events prompts several complications as well. During periods of heightened activity, it is possible for a single Telecom event to be embedded among several pages of another subsystem's events in the ISOE. It is easy to miss events, and sometimes the ISOE is so extensive that operators do not even attempt to scan it. Rather, they rely on an unofficial graphical sequence hardcopy product, the Spacecraft Flight Operations Schedule (SFOS), to monitor critical events. The SFOS, which is manually highlighted with a marker to indicate changes, creates problems when users unknowingly do not reference the latest activity modifications.

The current Voyager data display system presents another area of limitation. It allows only five plot display pages for the entire spacecraft team. The Telecommunications subsystem has control of a single page. One display page is capable of showing up to three plotted channels. In order to

change the plot parameters to select different channels to display, the operator must punch a card and feed it into the system's card reader. To obtain an additional plot, special permission must be secured from personnel of another subsystem who are willing to temporarily give up one of their own plots.

Broadened alarm limits present obvious complications. If, in fact, a component is in alarm within the broadened range, this condition will go undetected. For spacecraft activities which warrant an alarm limit change, and the operator chooses to forego the unwieldy paperwork process, then he must endure the false alarm for the remainder of that spacecraft activity.

A tabular display of spacecraft parameters is available which indicates alarms by reversing the color of the alarmed channel's field. However, this display is seldom used as the operator is usually viewing plotted data on the one allotted Telecom display page. As a result, a Telecom alarm condition generally is not detected by the Telecom operator until the Voyager Systems Analyst (who monitors and coordinates all subsystems) calls it to his attention.

Diagnosis

When a spacecraft or DSS parameter goes into alarm, the cause must be determined. In many cases, the condition is actually a false alarm due to inaccuracies precipitated by the limitations of the system. In other instances, the alarm exists because of common problems that occur on a frequent basis. For actual spacecraft problems, such as the failed radio receiver, hundreds of people must be notified and put on alert to solve the emergency. Regardless of the cause or severity of an alarm, a standard set of rules is routinely followed to determine the basis of the problem. Unfortunately, this knowledge resides with a select few, and the first rule of the standard

procedure is to consult the expert, even when the situation arises from a known false alarm.

THE SHARP SOLUTION

SHARP introduces automation technologies to the spacecraft monitoring process to eliminate much of the mundane processing and tedious analysis. The SHARP system features on-line data acquisition of all required information for monitoring the spacecraft and diagnosing anomalies. The data is centralized into one workstation and serves as a single access point for the aforementioned data as well as for the diagnostic heuristics. Figure 1 illustrates a top-level view of the SHARP system. Shown are the individual modules that comprise the system, as well as relevant components which are external to SHARP.

SHARP is implemented in CommonLISP on a SYMBOLICS 3650 color LISP machine. Many components of the system utilize STAR*TOOL [1] (patent pending), a language and environment developed at JPL which provides a tool box of state-of-the-art techniques commonly required for building AI systems. The SHARP system currently consists of approximately 40,000 lines of CommonLisp code, and STAR*TOOL comprises an additional estimated 85,000 lines of CommonLisp code.

Conventional Automation

The SHARP system captures raw predicts for on-line storage and processing. When the predicts are generated for the Voyager Spacecraft Team as hardcopy, the information is transferred over an RS-232C serial line to the SHARP system. Pass predicts may then be automatically generated at 15 second intervals, the shortest possible time interval between the arrival of any two spacecraft data points. Instantaneous predicts, which are pass predicts corrected in real-time for spacecraft pointing loss and DSS system

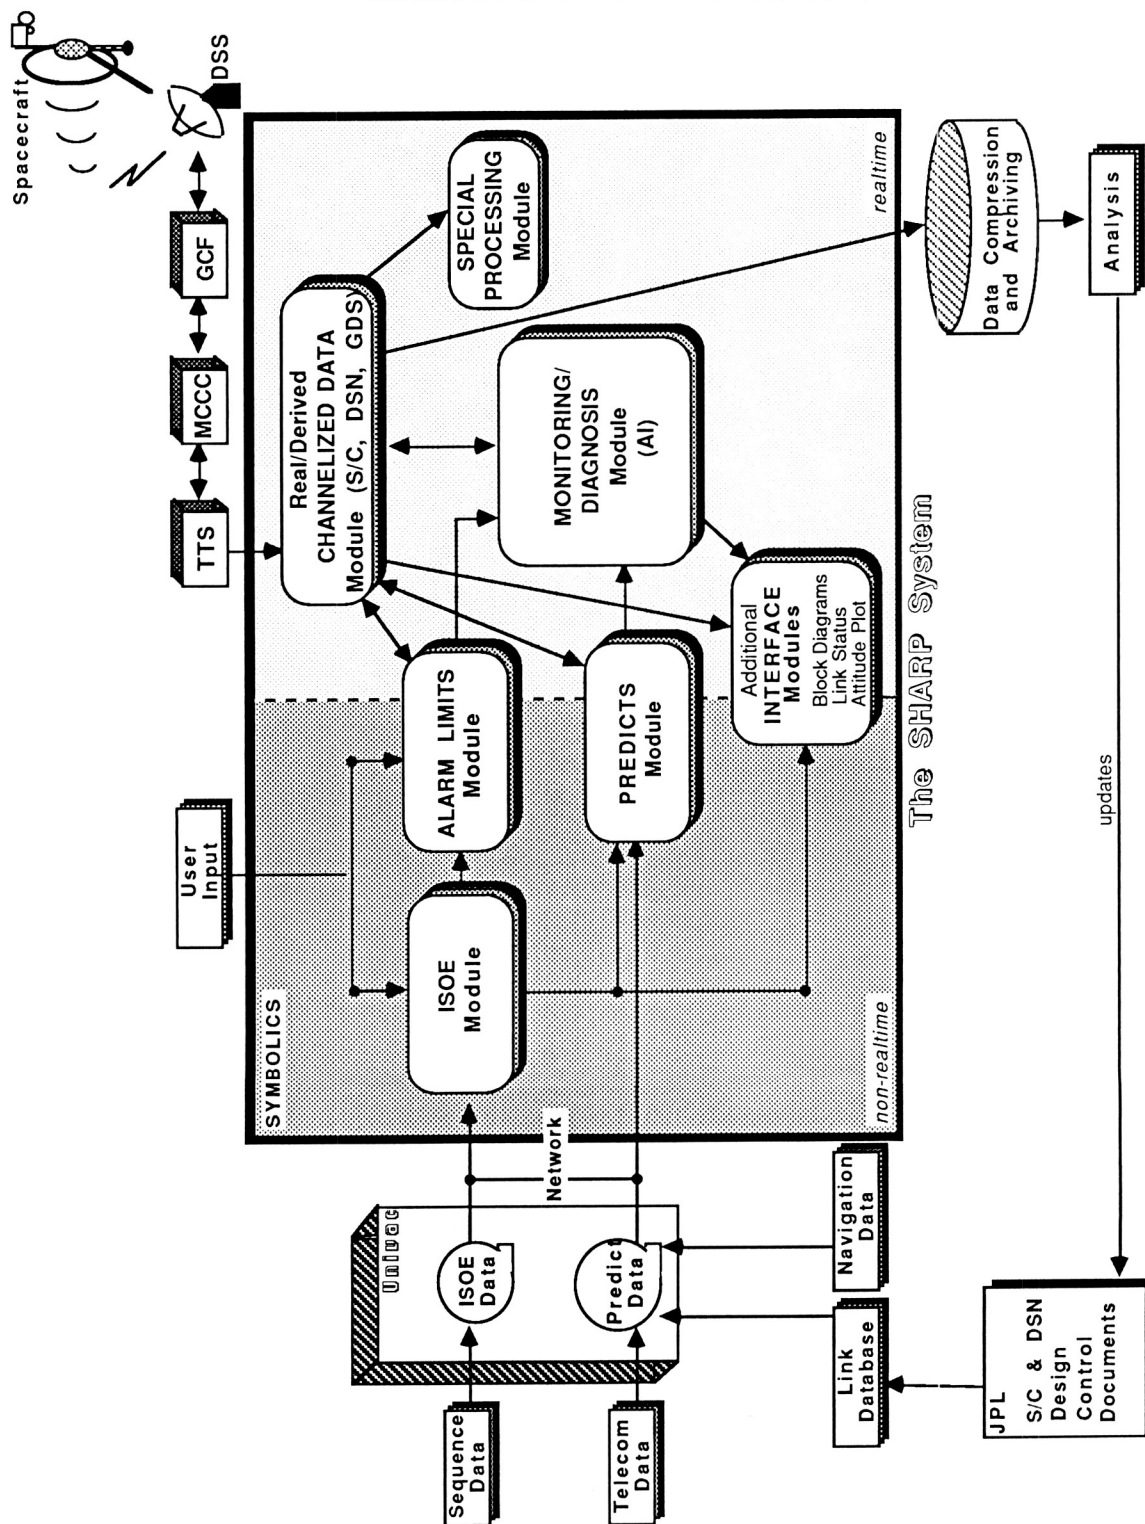


Figure 1. SHARP Telecom System Overview

noise temperature, are also automatically calculated at 15 second intervals. Spacecraft and DSS residuals, difference measurements between the actual values and predicted values, are automatically derived in real-time.

SHARP also acquires the Integrated Sequence Of Events for on-line storage and viewing. A generic capability to extract subsystem specific information has been developed, hence Telecom-specific events may be stripped from the ISOE and displayed to enable rapid identification of significant Telecom activities to be monitored during any particular pass. Editing capabilities facilitate on-line additions, deletions, and other changes to the ISOE, thus reducing the likelihood of referencing outdated material.

New plotting capabilities for the channelized data have also been implemented within the SHARP system. The operator can construct as many data plots per page, or screen, as desired, although five plots per page seems to be the optimal number for effective viewing. The user also possesses the capability to construct multiple pages which can be set as a program parameter. The user can change the display of plots on any given page at any time with simple menu-driven commands.

Alarm tables have also been constructed as part of the conventional automation process, and placed on-line within the SHARP system. A table for each relevant spacecraft configuration exists, resulting in alarm limits representative of the true thresholds for each data channel. SHARP determines alarm limits dynamically in real-time and accurately reflects each spacecraft or DSS configuration change. Dynamic alarm limit determination eliminates the cumbersome alarm change paperwork process as well as many occurrences of false alarms.

Several graphical displays in SHARP

automatically highlight alarmed events as they occur. These displays offer information ranging from the location of a problem to the probable cause of the alarm.

Enhanced Graphical Capabilities

The SHARP system provides numerous sophisticated graphical displays for spacecraft and station monitoring. A comprehensive user interface has been developed to facilitate rapid, easy access to all pertinent data and analysis. Displays have been constructed which range from placing data on-line to the creation of detailed graphics which provide a multitude of information at one glance. An interface exists for each major module of the SHARP system. Each interface provides customized functions which allow data specific to that module to be easily accessed, viewed, and manipulated. Each SHARP module can be accessed from any other module at any time, and all displays are in color with mouse sensitivity and menu-driven commands. Figure 2 shows the SHARP top-level system status view.

The Predicts interface in SHARP allows tabular display of raw predicts, pass predicts, instantaneous predicts, and residuals for any specified time range. A color-coded DSS availability graph has also been designed which enables rapid identification of available stations for any given viewing period. Situations which mandate that another Deep Space Station be acquired can hence be addressed immediately as opposed to the more arduous current method which requires the manual look up of each station at the specified time period.

The SHARP system provides an Integrated Sequence of Events interface which offers numerous capabilities to the operator. On-line viewing of any ISOE is available, and intricate modifications may be performed with ease. Editing the ISOE is accomplished via menu-driven commands which contain explanations of the complex ISOE data. For example, CC3A32330 means that the




Figure 2. SHARP Top Level System Status View.

x-band modulation index is 32, the two drivers are on, the subcarrier frequency is high, and the data line rate is high. Translation of these spacecraft commands from their raw form into more understandable summaries of spacecraft activity may be performed, and the user can request status summaries of any activity. A history display is maintained as the ISOE is updated so that the user can verify modifications.

SHARP's display which plots the channelized data, illustrated in Figure 3, is a significant improvement over existing capabilities. The user dynamically customizes the display at any time by selecting which and how many channels to view, the time scale, the data range for each plot, and even the icon to use for graphing points on each channel. Each plot is color-coded by the user for easy visual distinction between displayed channels. When any channel is in alarm, its corresponding data points are plotted in red, facilitating rapid detection of an alarm condition. The channel's associated alarm limits may be optionally overlaid onto the channel's plot for further information. Each data point is mouse sensitive to provide time and numerical value indicators, and an automatic counter continually indicates the number of data points per plot. Pan and zoom features augment this display, which can represent information as graphs of actual or derived data vs. time, xy plots, scatter plots, or logarithmic scales.

The SHARP system also provides an alarm limit interface which allows on-line viewing and editing of established spacecraft engineering alarm limits, DSS performance limits, ground data system limits, and residual thresholds. Authorized users may permanently alter any of these limits, and specified values may be changed temporarily for the remainder of that particular spacecraft pass. The later capability, manual override, enables alarm suppression or closer scrutiny for

any particular event with no intervening paperwork.

Among the new graphical analysis capabilities provided by the SHARP system is the Telecom link status display, as shown in Figure 4. Actual station coverage is illustrated, along with spacecraft transmitter power status, data rate, data outages, and real-time recording of station uplink (signal transmission) and projected downlink a round trip light time later. Detailed analysis is performed and information is subsequently color-coded to represent changes in status. The display provides the user with such valuable information as time ranges and explanations of data outages (i.e. no station coverage or ongoing spacecraft maneuver), and can warn the operator when to expect noisy data and why.

The SHARP system also features on-line functional block diagram schematics of the end-to-end communications path from the spacecraft through the Deep Space Communications Complex and Ground Communications Facility to the Mission Control Center at JPL and final destination of the Test and Telemetry System computers. Each top level system status may be viewed at successive levels of detail. The Telecom subsystem is very comprehensive, as spacecraft schematics have been developed for the all of its individual components. These dynamic block diagrams are driven by various ISOE status indicators and the channelized data. The status of spacecraft and DSS components (operational, off-line, or in alarm) is depicted by color, facilitating rapid status identification at a glance. Figure 5 shows the Telecommunications subsystem and Figure 6 illustrates the spacecraft receiver with associated diagnostic messages.

Another graphical display which combines various sources of information and data is SHARP's Attitude and Articulation Control display. This display




Figure 3. Plotting capabilities available in SHARP.




Figure 4. SHARP Telecommunications Link Status display.




Figure 5. SHARP schematic block diagram of Voyager Telecommunications subsystem.

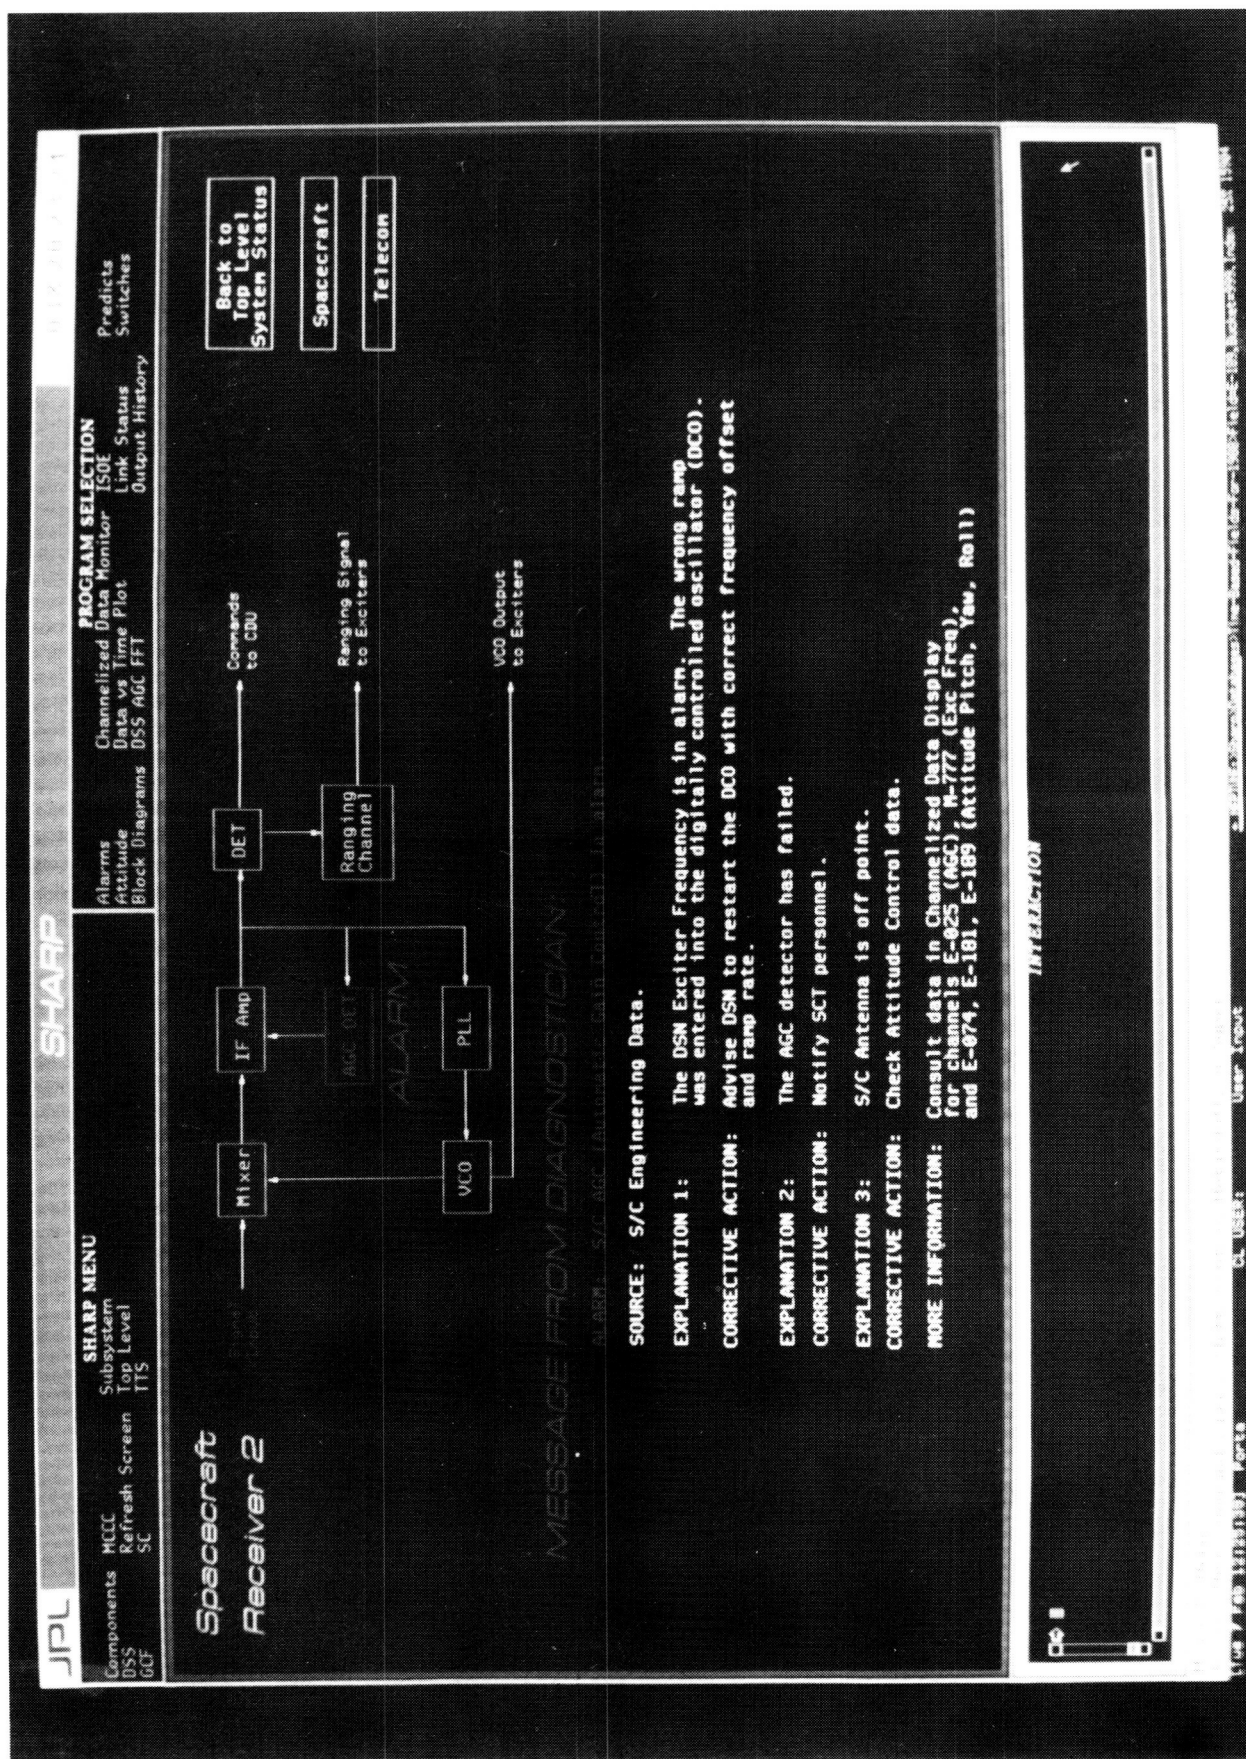


Figure 6. SHARP schematic block diagram of spacecraft receiver with message from diagnostician.

combines spacecraft motion parameters (pitch, yaw, and roll) and projects spacecraft movement over time. A limit cycle box which represents defined spacecraft deadband limits encloses the spacecraft icon. Alarm conditions are easily detected as the spacecraft icon drifts outside of the designated deadband box. Trailing vectors establish the path of the spacecraft.

The SHARP system also contains special processing modules to perform subsystem specific analysis. For the Telecommunications subsystem, a Fast Transform (FFT) of the DSS conical scanning component is performed to indicate when the antenna is going off point. This is a relatively common event, which currently may take hours to detect and correct. Spacecraft and scientific information can be permanently lost when this situation occurs. SHARP's FFT display illustrates the results of a Fast Fourier Transform process performed on 64 data points of a particular channel, and provides instant information on conical scan error. The problem can be detected in a matter of minutes, and the station can be contacted to correct the antenna movement prior to the loss of data.

Artificial Intelligence

The artificial intelligence modules of SHARP are written in an expert system building language called STAR*TOOL. STAR*TOOL is a programming language designed at JPL to meet many of NASA's demanding and rigorous AI goals for current and future projects. Appendix A contains a more detailed description of the STAR*TOOL system.

Artificial intelligence techniques are distributed throughout all components of the SHARP system. Intelligent programming methodologies such as heuristic adaptive parsing, truth maintenance, and expert system technology enable more effective automation and thorough analysis for SHARP functions. Fault detection becomes almost immediate with a greater

degree of accuracy and precision, and the system quickly generates fault hypotheses.

A blackboard architecture, provided by STAR*TOOL, serves as a uniform framework for communication within the heterogeneous multi-process environment in which SHARP operates. Generally, when two or more processes are cooperating, they must interact in a more complicated manner than simply setting global variables and passing information along such paths. SHARP provides a standardized method of communication between multiple processes, which include real-time posting of incoming telemetry data and the monitoring of data networks.

Heuristic adaptive parsing is implemented for SHARP's raw predicts database. Periodically the format of this data source changes without mission operations being notified. Generally this would require the raw predicts parser to be rewritten to incorporate the new format. However, SHARP utilizes Augmented Transition Network [2] (ATN) techniques to accomplish adaptive parsing. The advantage of such an ATN lies in its ability to parse the database according to semantic content rather than syntactic structure. The raw predicts database can therefore be modified and yet remain successfully parsable. This heuristically controlled format insensitive parsing ensures continuity despite format modifications in predict generation.

The centralized database of the SHARP system serves as a central repository of all real-time and non-real-time data, and functions as a local buffer to enable rapid data access for real-time processing. Numerous database manipulation functions have been implemented, and database daemons have been constructed to implement spontaneous computations. Requests can be made to the database to trigger arbitrary activities when a complex combination of past, present, and future

events occur. A wide selection of retrieval methods by time or value highlight the flexibility inherent in the database. Requests to the database can be made from both AI and non-AI modules of SHARP, and can be handled serially or in parallel.

Various SHARP modules represent and manipulate data symbolically rather than numerically so that particular numeric values can change without forcing the algorithms themselves to be modified. For example, to determine if a channel is in alarm, the rule interpreter manipulates one symbolic fact, *ChannelInAlarm*, rather than the many numeric operations that are required to make an actual determination. This is a significant advantage as SHARP presently analyzes over 100 channels, and the alarm determination process varies from channel to channel. Symbolic representation and manipulation of data also simplifies the exchange of information between SHARP modules and reduces reliance on specific dimensionless numeric values.

The diagnostic component of SHARP is composed of a hierarchical executive diagnostician coupled with cooperating and non-cooperating mini-experts. Each mini-expert is responsible for the local diagnosis of a specific fault or class of faults, such as particular channels in alarm, conical scan errors, or loss of telemetry. A non-cooperating expert focuses only on its designated fault area, but a cooperating expert has the additional capability of searching beyond its local area to identify related faults that are likely to occur. Cooperating experts are used in situations where the identification of a particular fault cannot be made by examining a single fault class alone.

The executive diagnostician combines input propagated from each local diagnostician and reviews the overall situation to propose one or more fault hypotheses and recommended corrective actions. When multiple fault hypotheses

are generated, the system lists all possible causes of the anomaly and ranks each according to plausibility.

If one or more of the cooperating experts fails, the executive diagnostician will continue to operate with only a reduction in the area of local diagnosis that would have been derived from the failed mini-experts. Similarly, if the executive diagnostician fails, the cooperating experts will locally diagnose the faults in isolation of multiple fault consideration.

The diagnostician is implemented in rules that execute in pseudo-parallel in pursuit of multiple hypotheses. Pseudo-parallelism is implemented in SHARP using facilities provided by STAR*TOOL, which includes parallelism as a fundamental control structure. The diagnostic rules operate in isolation of one another by executing in independent contexts provided by the STAR*TOOL memory model, and communicate through the Blackboard facility.

These contexts can be organized into a tree-like structure to represent contradictory information resulting from changes in facts or from the introduction of new or contradictory hypotheses. Facilities in the truth maintenance system handle data and demand driven diagnoses to ensure an appropriate balance between the persistence of hypotheses and sensitivity to new data.

Bayesian inference processes are used for comparing multiple hypotheses and for prioritizing conflicting fault hypotheses. Bayesian inference procedures also perform uncertainty management to allow continued high performance in the presence of noisy, faulted, or missing data.

The truth maintenance system constantly monitors for violations of logical consistency. For example, it performs conflict checking to maintain

consistency among multiple rule firings, hypotheses, and the knowledge base, and allows the context sensitive management of alarms through a complex response system to combinations of alarm conditions. Truth maintenance techniques also provide a variety of functions for temporal reasoning in multiple fault diagnosis.

CONCLUSIONS

Spacecraft and ground data systems operations present a rigorous environment in the area of monitoring and anomaly detection and diagnosis. With a number of planetary missions scheduled for the near future, the effort to staff and support these operations will present significant challenges.

The SHARP system is an attempt to address the challenges of a multi-mission monitoring and troubleshooting environment by augmenting conventional automation technologies with state-of-the-art artificial intelligence. Results of this effort to date have already shown significant improvements over current Voyager methodologies and have provided enhancements to several aspects of Voyager operations.

This type of automation technology will endow mission operations with considerable benefits. In as many areas as are automated, expert knowledge will be captured and permanently recorded, reducing the frenzied state that occurs when domain specialists announce their impending retirement. Cost reductions will occur as a result of automation and decreased requirement for 24-hour real-time operator coverage. Automatic fault detection and analysis will facilitate quicker response times to mission anomalies and more accurate conclusions. The time savings afforded by SHARP-like capabilities, especially during periods of unmanned operation or during emergencies, could mean the difference between the loss or retention

of critical data, or possibly even the spacecraft itself.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the following people for their participation in the SHARP project: David J. Atkinson, task management; Harry J. Porta and Gaius Martin, software development; Boyd Madsen, expert knowledge for Voyager and Galileo spacecraft telecommunications; and Bruce Elgin and Erann Gat, software contributions.

The work described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- [1] James, Mark L., *STAR*TOOL Reference Manual*, Internal Document, Jet Propulsion Laboratory, Pasadena, California (1988).
- [2] Charniak, Eugene and McDermott, Drew, *Introduction to Artificial Intelligence*, Addison-Wesley (1985).

APPENDIX A: STAR*TOOL

Knowledge-based systems for automated task planning, monitoring, diagnosis, and other applications require a variety of software modules based on artificial intelligence concepts and advanced programming techniques. The design and implementation of such modules requires considerable programming talent and time, and a background in theoretical artificial intelligence. Sophisticated software development tools that can speed the research and development of new artificial intelligence applications are therefore highly desirable. The STAR*TOOL system, designed and built by

Mark James, was developed specifically for this purpose. Included in the system are facilities for developing reasoning processes, memory-data structures and knowledge bases, blackboard systems, and spontaneous computation daemons.

Computational efficiency and high performance are especially critical in artificial intelligence software. This consideration has been an important objective of STAR*TOOL, and has led to its design as a toolbox of AI facilities that may be used independently or collectively in the development of knowledge-based systems.

STAR*TOOL provides a variety of facilities for the development of software modules in knowledge-based reasoning engines. The STAR*TOOL system may be used to develop artificial intelligence applications as well as specialized tools for research efforts.

STAR*TOOL facilities are invoked directly by the programmer in the Common LISP language. For improved efficiency, an optional optimization compiler was developed to generate highly optimized CommonLISP code.

STAR*TOOL was designed to be efficient enough to operate in a real-time environment and to be utilized by non-LISP applications written in conventional programming languages such as ADA, C, Fortran, and Pascal. These non-LISP applications can run in a distributed computing environment on remote computers, or on a computer that supports multiple programming languages.