A BEACON MONITORING SYSTEM FOR AUTOMAI MANAGEMENT OPERATIONS

Michael A. Swartwout and Christopher A. Kitts
Faculty Advisor - Professor Robert J. Twiggs

Space Systems Development Laboratory
Department of Aeronautics and Astronautics
Stanford University, Stanford, California 94305

ABSTRACT

Stanford's Space Systems Development Laboratory (SSDL) has initiated a new space system technology initiative in order to develop, demonstrate, and validate a beacon monitoring system for spacecraft. This system consists of automated fault detection on board a spacecraft, a state of health beacon signal broadcast by the spacecraft, a ground based monitoring network, and a mission control center capable of efficiently integrating this health assessment strategy into its operating architecture. SSDL is investigating this technique by identifying fundamental design drivers, developing a system responsive to these drivers, and deploying the resulting system on microspacecraft and within SSDL's developing, global, automated space operations network. This paper reviews the beacon monitoring concept, describes the design criteria for such an operations strategy, and presents the current development of the SSDL beacon system.

INTRODUCTION

The spacecraft industry regularly cites its failure to incorporate operational concerns into system design drivers. This fault typically leads to inefficient and human intensive mission control activities that can account for up to 60% of total program costs.

Autonomy is often praised as a technology capable of reducing costs and enhancing performance in space mission operations. Through careful deployment within the overall mission architecture, automation can augment or replace human decision making in order to increase reaction speeds, reduce errors and stress, mitigate cognitive overload, enhance safety, lower costs, focus analysis, cut bandwidth requirements, and free human reasoning for strategic tasks requiring high levels of robustness.

System health management is a specific mission operations task that has been enhanced through the use of automation technologies. The space industry routinely uses ground based expert systems to analyze spacecraft telemetry and detect the existence of faults. While this model has proven to be beneficial in reducing the workload of human controllers during nominal operations, it still requires full use of scarce ground equipment and bandwidth resources in order to deliver spacecraft telemetry to the control center.

To address this drawback, many developers in the spacecraft community advocate the migration of detection capability from the ground to on board the spacecraft. This new model requires the on-board reasoning system to perform realtime health assessment and to use a "beacon" to report the vehicle's status to the mission control center. Composed of, at most, a few bits of information, this beacon signal will summarize the spacecraft's status. When healthy, a "Normal" or "I'm OK" signal will reduce the need for routine health assessment contacts thereby conserving resources. When a fault exists, an "Emergency" or "Help Me" signal will trigger notification of controllers and can be used to initiate a variety of contingency operation functions.

The Air Force is studying the feasibility of such a system for incorporation into the Air Force Satellite Control Network (AFSCN).3 Similarly, NASA's Jet Propulsion Laboratory is planning on implementing such a system within the Deep Space Network for its Pluto Express mission as well as its series of New Millennium Program spacecraft.4 As these organizations are learning, however, the beacon monitoring model of fault detection poses a number of additional design challenges that must be considered. Because of the nature of these networks, however, implementation and validation of such systems has been slow. To prove the value of the beacon monitoring concept, SSDL is preparing to conduct a real world evaluation through the use of microspacecraft and a global, automated space operations network.

DESIGN DRIVERS FOR A BEACON MONITORING SYSTEM

As has been stated, the ongoing industry and government programs fordesigning beacon monitoring systems clearly point to a variety of benefits in reducing human and ground resource requirements for nominal spacecraft operations. But because this model radically alters the location of decisionmaking authority and the distribution of information, beacon monitoring also creates a variety of additional design challenges that must be addressed by the overall space system.

First, the on-board monitoring process must be robust enough to detect complex and unanticipated faults while also reducing false alarms. While the current crop of ground based expert systems can be fragile in this respect, the availability of the full spacecraft telemetry stream at the mission control center permits simple and timely human intervention when required. Since beacon monitoring reduces the amount of telemetry available for human analysis, advances in detection capability are required.

Second, the beacon signal itself must convey a small but appropriate amount of information in order to initiate actions throughout the ground segment. The choice of the number of bits in the signal is therefore a direct function of the responsiveness of the mission architecture. Also, the beacon signal should be broadcast on a nearly continuous basis in order to provide quick notification to the control center. Third, a suitable beacon receiving network is required. Typical spacecraft operations require directional and/or dedicated ground station equipment in order to contact vehicles. The whole notion of beacon monitoring, however, is based upon conserving resources. For this reason, a communications link based on exploiting excess station capacity or simultaneously utilizing equipment is required. The previously stated desire for timely signal reception further leads to a requirement for deploying beacon receiving equipment throughout the world. Of course, these distributed receiving stations must be linked to the mission control center through some sort of readily available communications network.

Fourth, the mission control center must efficiently exploit a received beacon signal by automatically notifying operators, allocating resources for timely contact with the vehicle, and preparing engineers for in-depth vehicle analysis. Concerning this last point, current use of automated ground based fault detection systems is already highlighting the challenge of keeping human operators fully apprised of the spacecraft's state and operating conditions. Because beacon monitoring can be used to reduce the amount telemetry being transmitted to the ground, this challenge is magnified. Additionally, a sudden Emergency signal, especially after months of nominal and unattended operation, can generate significant operator stress which can lead to poor judgment and errors. This can be partially addressed through automated referral to contingency action plans and engineering documentation.

THE SSDL TESTBED SYSTEMS

The beacon monitoring investigation is drawing upon two testbed systems already under development within SSDL. The first of these is the Stanford Audio Phonic Photographic Infrared Experiment (SAPPHIRE) microspacecraft, the first satellite of SSDL's Satellite Quick Research Testbed (SQUIRT) program.5 The second is a new operations research project known as the Automated Space System Experimental Testbed (ASSET) program.

The SAPPHIRE Microsatellite[6]

The SAPPHIRE satellite, shown in Figure 1, has three primary missions. The first of these is the space characterization of newly developed micromachined tunneling infrared sensors. These non-cryogenic sensors have been jointly developed by the Jet Propulsion Laboratory and Stanford University. The other missions are providing photographic and voice broadcast services to the public through the use of modified, commercially available, non-space rated components.

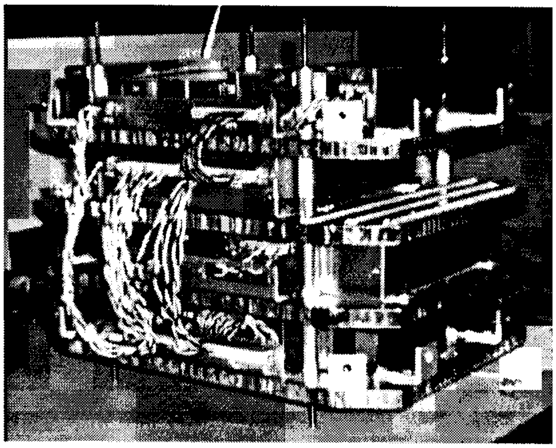

Figure 1 - SAPPHIRE Bus in Integration Testing

The SAPPHIRE bus is a 35 pound hexagonal cylinder measuring 8" on a side and 11" tall. The structure is constructed primarily from 0.5" aluminum honeycomb and consists of four modular, stacked subsystem trays.

The power tray contains a single ten cell NiCad battery and 5/12V DC-DC converters. The communications tray consists of a modified HAM radio transmitter, receiver and terminal node controller. The processing tray holds the Motorola 68332 CPU, 256 Kb of RAM, 256 Kb of hardened ROM, and serial interfaces. The payload tray contains two banks of IR sensors, a modified Logitech Fotoman camera, and a modified voice synthesizer. SAPPHIRE is finalizing its operational testing and is awaiting integration with a launch provider.

The ASSET Program[2]

The ASSET system is a global space operations network under development within SSDL. This first goal of this system is to enable low-cost and highly accessible mission operations for SQUIRT icrosatellites as well as other university and amateur spacecraft. The second goal of this system is to serve as a comprehensive, low inertia, flexible, real world validation testbed for new automated operations technologies. Figure 2 shows a high level view of the ASSET mission architecture. The basic components include the user interface, a control center, ground stations, communications links, and the target spacecraft. During the current development phase, a highly centralized operations strategy is being pursued with nearly all mission management decision making executed in the control center. These tasks include experimental specification, resource allocation throughout the ground and space segment, fault management, contact planning, data formatting and distribution, and executive control.

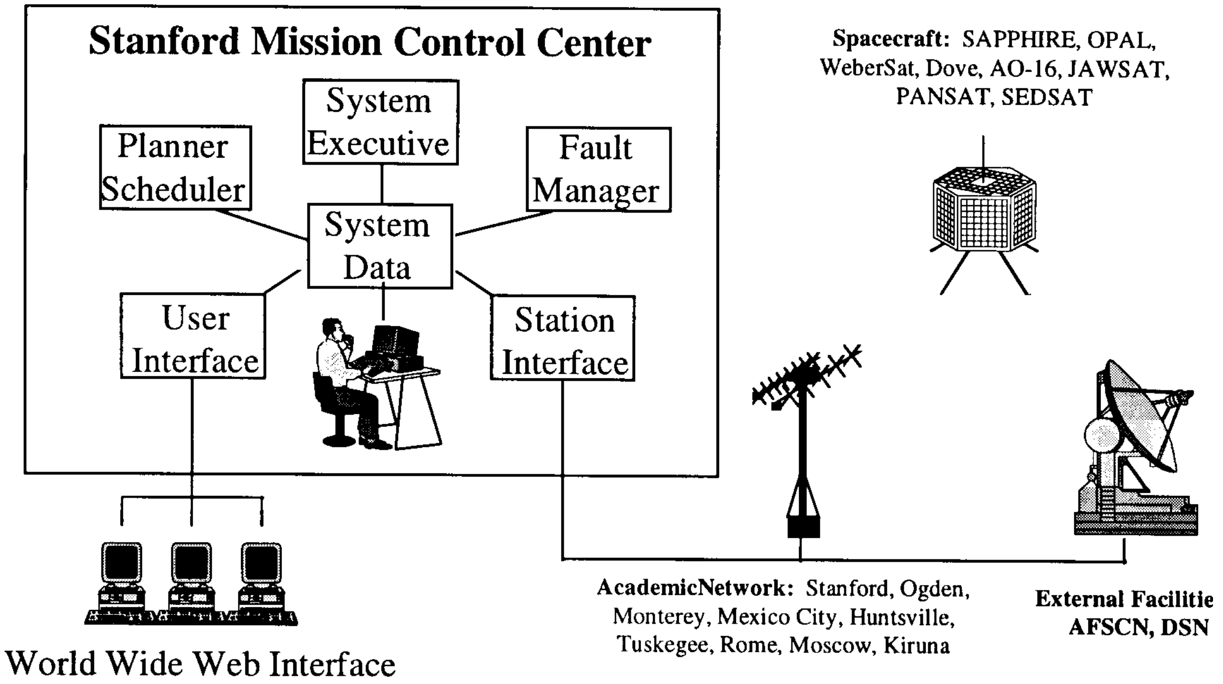

Figure 2 - The ASSET Space System Architecture

WeberSat, a Weber State University and Amateur Satellite Organization spacecraft launched in 1990, and SAPPHIRE, the first SSDL SQUIRT spacecraft currently being completed, are presently being integrated into the ASSET system. In addition, controllers for five other satellites have expressed interest in becoming part of the system. As for ground stations, the Weber State and Stanford ground stations are the first two facilities to be included. Seven other stations throughout North America and Europe have been identified for future integration.

IMPLEMENTATION OF THE SAPPHIRE/ASSET BEACON SYSTEM

While general design drivers for an optimal beacon monitoring system have been outlined, the specific implementation of such a system within the SAPPHIRE/ASSET testbeds is limited due to computational capability and development resources. While this is unfortunate, it is believed that a partial implementation of an ideal system will be enough to judge the value of the beacon strategy.

This section outlines the design and implementation of the system currently under development. In doing so, it describes the on-board capabilities being added to SAPPHIRE, the beacon receiving system, and the response actions being integrated into ASSET.

SAPPHIRE Beacon Generation

Functional capability on the SAPPHIRE vehicle is governed by the spacecraft being in either "Standard Mode" or "Safe Mode." SAPPHIRE usually operates in Standard Mode. In this mode, all satellite functions, such as payload operations and multi-user contacts, are fully available given the constraints of standard operating procedures. Safe Mode, on the other hand, terminates previously configured payload activity, places the spacecraft into a low power mode, and limits spacecraft functionality by permitting only a single "superuser" to command the spacecraft. Safe Mode is automatically initiated on-board SAPPHIRE based upon predetermined critical conditions such as a low battery voltage. Safe Mode is also the default CPU mode so that unanticipated CPU rebooting will limit the vehicle's functionality until the spacecraft is properly configured.

Given these two levels of vehicle functionality, the first question for the design of the beacon monitoring system concerned the resolution of the signal itself. Although the spacecraft modes suggest the use of a two state beacon, the use of this criteria for selecting the number of beacon states is inappropriate. The beacon function is essentially a transform that maps the large set of spacecraft states to a concise set of appropriate ground system reactions. So while various spacecraft states are reduced to one of two functional levels, a broader set of ground reactions is relevant.

For example, while the vehicle senses and safes itself due to critical conditions, routine telemetry analysis can often permit forecasting of future emergencies so that anomalies can be more gracefully addressed or even averted. SAPPHIRE has basic on-board telemetry limit checking capability which

Mode	Beacon Mode	Vehicle Mode	Vehicle Status
1	Normal	Standard	Healthy
2	Abnormal	Standard	Out of Limit Telemetry
3	Critical	Safe	CPU Controlled Safe Mode
4	Emergency	Safe	CPU Reset Induced Safe Mode

***Table 1** - SAPPHIRE Beacon and Vehicle Modes*

can be used to signal operator attention to potential faults. In addition, while the vehicle's Safe Mode will always limit the level of functionality on-board SAPPHIRE, the means of entering Safe Mode has a direct consequence on how the mission control center should recover operations. For example, Safe Mode due to a CPU reset will have cleared the spacecraft's time, RAM based operating biases, mission data, and stored commands. Less drastic events occur during a CPU controlled Safe Mode. For these reasons, a four tone beacon signal is being implemented as shown in Table 1.

The SAPPHIRE CPU controls the beacon state based on its understanding of realtime telemetry. This information is compared with a table of upper and lower limiting values for important signals such as component temperatures, currents, and voltages. The limit values are commandable so that analysis can evolve with normal operating conditions. For example, solar panel current limits will be modified over time so that solar cell degradation will not be detected as a fault. In addition, the SAPPHIRE limit checking system can operate at an aggregate level in order to provide more focused analysis. For instance, the current from a single solar panel will typically vary between zero and some maximum normal level of output. An additional and much stronger check can be made, however, by ensuring that all solar cells are not near their maximum allowable level at the same time (SAPPHIRE's geometry prevents the sun from illuminating all panels at the same time). Finally, SAPPHIRE's limit checking is context sensitive so that different expected ranges are applied to realtime telemetry based upon the estimated state of the vehicle. An example of this is that during an eclipse the battery should be discharging. Satisfaction of all of these limit checking routines will maintain the beacon in its Normal Mode setting (Mode I). Violation of any of these checks will trigger Abnormal Mode (Mode 2).

While Abnormal Mode serves as an indicator for ground operators, violation of a few specific telemetry checks will result in a CPU initiated Safe Mode. In this mode, all instruments are shut off, the command database is cleared, the spacecraft is placed into its low power mode, and CPU access is limited to a single superuser. In addition, the beacon system is placed into Critical Mode (Mode 3). Battery voltage is the primary check for initiating this sequence of activities. A significant drop in battery voltage means that it is overloaded; persistent use of the spacecraft would jeopardize battery survival and thus the mission. It takes operator instructions to bring SAPPHIRE out of any Safe Mode; thus, even a return to normal telemetry ranges will not result in a return to Standard Mode.

As defined in SAPPHIRE's boot code, a CPU reset will place the spacecraft in vehicle Safe Mode and in Emergency beacon mode (Mode 4). Such an event is caused by gross power mismanagement, watchdog timer resets, or ground command; as such, its occurrence is more serious than a CPU initiated Safe Mode and needs to be specially treated. Operator commands are required to return to Standard Mode.

Figure 3 depicts a high level diagram of SAPPHIRE's on-board fault handling. As has been seen, SAPPHIRE relies on analysis of its telemetry values to identify problems. Obviously, faulty telemetry lines are a problem that can disguise faults or generate false alarms. If a bad telemetry line is identified, it can be effectively removed from the limit tables. In addition, the entire beacon monitoring process is a commandable option; should it be determined that the CPU can no longer manage this system, it can be turned off. Also, vehicle Safe Mode is an end state; once SAPPHIRE enters this mode, it no longer takes any action based on the telemetry checks. Finally, changes in vehicle and beacon modes are logged in an activity file on the vehicle.

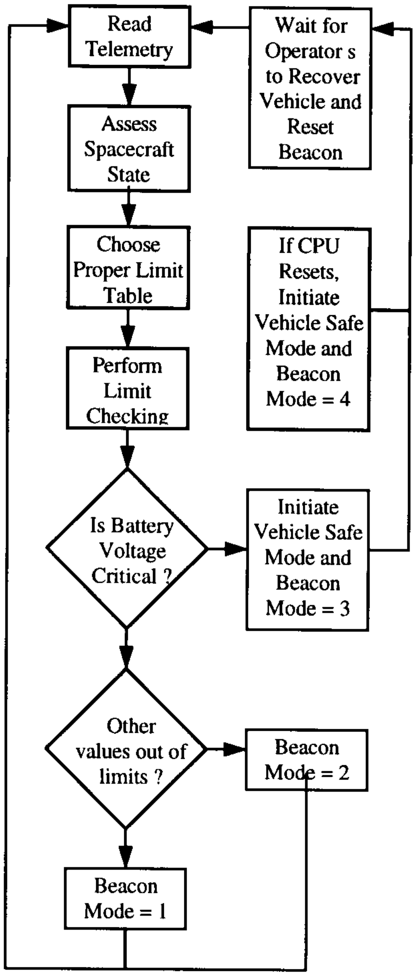

Figure 3 - SAPPHIRE's On-Boar

Fault Handling

Beacon Reception

The next stage of the monitoring system is the reception of the beacon signal. Unless SAPPHIRE is in contact with ground operators, the beacon is broadcasting at all times. Simple, receive-only, HAM radio frequency stations are being designed to acquire the two-bit beacon signal. These stations will accommodate multiple vehicles and will forward received signals to mission control centers through Internet or modem connections. The development of these stations is being conducted through educational initiatives with Tuskegee University and the Space Engineering School in Kiruna, Sweden. For test purposes, SSDL is using small, portable, omnidirectional, laptop based receiving stations in order to develop the end-to-end beacon monitoring system. Because the most effective and responsive use of the beacon requires a high degree of signal availability, these receive stations are planned for deployment at every ground station currently being integrated into the ASSET system. This will provide a high level of coverage across the United States and Europe. An initiative with the Air Force is currently underway to locate some of these systems at various AFSCN ground stations throughout the world. This will significantly enhance coverage of the beacon network.

Beacon Response

Based upon the state of the beacon signal, the mission control center must perform a set of appropriate actions. The beacon's value is always recorded. If the mode is anything other than Normal, an automated notification function informs operators of the beacon's state. This can be achieved through electronic mail, audible alarms, and automated paging services. In Abnormal mode, the scheduling of additional contact with the vehicle is currently being left to the discretion of the operator on call.

If the beacon is in either Critical or Emergency mode, a number of more involved functions are performed. First, SAPPHIRE experiments are postponed or canceled. Second, ground station support for SAPPHIRE is scheduled. Third, the schedules of other ASSET spacecraft are reworked to account for changes in ground station availability and experimental loading. While these functions are currently manual, automated scheduling for multiple spacecraft and multiple ground stations is being developed.

Although not formally part of the initial SAPPHIRE/ASSET beacon system, automated documentation retrieval and advanced fault isolation, diagnosis, and recovery techniques are being developed in order to aid the operator once full fledged contingency analysis begins. This capability is essential to the success of the beacon concept since operators will not be intimately familiar with spacecraft behavior once automated fault detection capability is established and trusted. In later versions of the beacon system, onboard fault summarization will be included; this will enable the advanced documentation and fault management functions to be integrated directly into the overall beacon monitoring system. An open issue in the design of ASSET's response is how to deal with on-board mission data when the CPU initiates a vehicle Safe Mode. Clearly, after a CPU reset, all data is lost, and ASSET must terminate active experiments or reschedule them to collect data that was lost. But in a CPU initiated Safe Mode, the data is likely to still be on board and intact. Because of fault analysis and recovery requirements, there may be some time before this data can be returned to the control center, or it may be lost during the recovery process. To ignore data that is available, however, is an inefficient use of limited spacecraft and ground resources. The trade, then, is to determine the conditions under which the data "trapped" on Sapphire should be salvaged or given up for lost. This issue is being resolved through experimentation and assessment of the beacon system.

FUTURE WORK AND INITIATIVES

The beacon monitoring system is currently being implemented and tested in the SAPPHIRE and ASSET systems. Since SAPPHIRE has not yet been launched, the designers can take advantage of this situation by heating, cooling, or disconnecting various components while other operators try to identify and correct for such faults without knowing what exactly has been done. The automated system will be compared to human-based reasoning in order to assess the capability of the system.

Improvements in automated fault management are being pursued and, where applicable, will be integrated into the SAPPHIRE/ASSET system. This work includes refinement of anomaly detection techniques and the system level corroboration of suspect telemetry readings. Advanced work in expert systems and model based reasoning will also be applied to automated diagnosis and fault recovery operations.

The current ASSET scheduling system relies on a simple first come first serve strategy. Look ahead scheduling techniques will soon be instituted; this will permit a more efficient use of the overall space system based on user priority, experimental value, and available resources.

The operator interface currently consists of a World Wide Web based system to display and filter raw telemetry files. More mature interfaces are being developed in order to enhance health analysis. One particular project consists of displaying system status through the use of a three dimensional representation of the spacecraft. Using the Virtual Reality Modeling Language, a spacecraft model is being developed and can be viewed through the Web. This method displays temperatures as color, explicitly displays orbital position and vehicle attitude, and can be slaved to electronic documentation.

Finally, SSDL is actively pursuing initiatives with external organizations in order to expand its work on the beacon monitoring system and other projects. As has been mentioned, cooperative efforts with students at Tuskegee University and the Space Engineering School in Sweden are encouraging broa~ participation at all levels of design and implementation of the beacon network. Also, the Air Force Satellite Control Network has specified the SSDL operations group as a "Pathfinder" in its effort to reduce costs and commercialize the AFSCN. Finally, cooperation with NASA's Deep Space Network is ensuring that the beacon work is relevant and directly applicable to distributed mission architectures.

CONCLUSIONS

Beacon monitoring is a technique that combines on-board anomaly detection capability with a reactive ground system. This method provides automated fault detection while also saving large amounts of human and communications resources during nominal spacecraft operations.

This paper has discussed the design and development of a beacon monitoring prototype for spacecraft operating in a shared, global mission architecture. To begin, the fault detection system on SAPPHIRE, while applied to a simple spacecraft, consists of adjustable and context-sensitive limit checking; this capability is more advanced from a functional perspective than many complex spacecraft operating today. Next, SAPPHIRE's beacon resolution has been explicitly determined through an analysis of the appropriate levels of ground system response to various spacecraft states. For the ground segment, design drivers for the beacon receiving stations have been specified. The design and strategic installation of these stations is being coordinated through SSDL's academic partners. Finally, precise mission control responses have been defined and are being implemented as part of the ASSET program. These automated responses include notification, planning, scheduling, fault management, and document retrieval. Each of these capabilities is being validated to ensure safe, efficient, and cost-effective management of the SAPPHIRE mission.

Beacon monitoring shows great promise as a method to drastically reduce the costs of spacecraft operations. Without such innovations, many of the large scale and aggressive missions of the future will simply not be possible. NASA's Pluto Express project, with a spacecraft that is subjected to a decade of interplanetary cruise, considers the use of beacon monitoring to be "crucial to the success" of the mission.7 The NASA New Millennium Program also considers beacon technology as a principal strategy for the missions of the next century; a demonstration of this technology is currently being incorporated into the program's Deep Space 1 mission. In addition, the Air Force Phillips Laboratory advocates the development of beacon systems as a means to reduce routine operating support for health monitoring and to increase confidence in the use of intelligent systems for spacecraft operations.8

The SAPPHIRE/ASSET beacon prototype is providing valuable feedback to the designers of these missions. This experiment offers a functionally mature system to investigate the real-world issues involved in automated spacecraft health monitoring. The technical simplicity, coupled with ASSET's flexible architecture, enables rapid experimentation and provides an ideal research testbed for the spacecraft community.

ACKNOWLEDGMENTS

The authors wish to acknowledge the outstanding work and commitment of other members of both the SAPPHIRE and the ASSET design teams. In addition, special thanks is given to the SAPPHIRE operations team for their willingness to participate in, contribute to, and provide feedback concerning various aspects of the beacon monitoring system. Finally, appreciation is extended to NASA's Jet Propulsion Laboratory and Ames Research Center for their support and assistance with this research project.

REFERENCES

[I] Ely, Neal, and Thomas P. O'Brien, "Space Logistics and Reliability", in Space Mission Analysis and Design, ed. James R. Wertz and Wiley J. Larson, 633-656. London: Kluwer Academic Publishers, 1991 .

[2] Kitts, Christopher A. "A Global Spacecraft Control Network for Spacecraft Autonomy Research." Pending publication in the Proceeding of SpaceOps '96: The Fourth International Symposium on Space Mission

Operations and Ground Data Systems in Munich. Germany. September 16-20. 1996, by the Electronic Publishing System. Munich, Germany: German Aerospace Research Establishment, 1996.

[3] Schultz,Michael. "Lifeline: A Concept for Automated Satellite Supervision." In the Proceedings of the Software Technology for Satellite Autonomy Workshop in Albuquerque, NM, June 22-25. 1993, by the Air Force Phillips Laboratory. Albuquerque, NM: Air Force Phillips Laboratory, 1993.

[4] Wyatt, E. Jay, and John B. Carraway, "Beacon Monitoring Approach to Spacecraft Mission Operations", NASA Jet Propulsion Laboratory White Paper, 1995.

[5] Kitts, Christopher A., and Robert J. Twiggs. "The Satellite Quick Research Testbed (SQUIRT) Program." In the Proceedings of the 8th Annual USU/ AIAA Conference on Small Satellites in Logan. UT, August 29 - September 1. 1994, by the Center for Space Engineering, Utah State University. Logan, UT: Utah State University, 1994.

[6] Lu., Richard A., Tanya A. Olsen, and Michael A. Swartwout. "Building 'Smaller, Cheaper, Faster' Satellites Within the Constraints of an Academic Environment." In the Proceedings of the 9th Annual USU/AIAA Conference on Small Satellites in Logan, UT. September 19-22, 1995, by the Center for Space Engineering, Utah State University. Logan, UT: Utah State University, 1995.

[7] Wyatt, E. Jay, Software Lead, NASA Pluto Express Mission. Interview by authors, 12 June 1996, Pasadena, CA

[8] Zetocha, Paul, Robert Statsinger, and David Frostman. "Towards Autonomous Space Systems." Air Force Phillips Laboratory White Paper, 1993.