AUTOMATED HEALTH OPERATIONS FOR THE SAPPHIRE SPACECRAFT

Michael A. Swartwout* **and Christopher A. Kitts**†
Stanford Space Systems Development Laboratory
Stanford, CA 94305-4035

ABSTRACT

Stanford's Space Systems Development Laboratory is developing methods for automated spacecraft health operations. Such operations greatly reduce the need for ground-space communication links and full-time operators. However, new questions emerge about how to supply operators with the spacecraft information that is no longer available. One solution is to introduce a low-bandwidth health beacon and to develop new approaches in on-board summarization of health data for telemetering. This paper reviews the development of beacon operations and data summary, describes the implementation of beacon-based health management on board SAPPHIRE, and explains the mission operations response to health emergencies. Additional information is provided on the role of SSDL's academic partners in developing a worldwide network of beacon receiving stations.

KEY WORDS

Health management, Beacon, Spacecraft operations, Data summary, Anomaly detection, Automation.

INTRODUCTION

Two trends are forcing significant changes in spacecraft operations. The first is pressure to reduce cost, such as ground-to-satellite communication links, operator man-hours, and operator training. The second is the growth of multi-satellite, coordinated missions; not only are there more spacecraft, but such missions introduce systematic problems related to constellation management. Taken together, these changes make the operations problem intractable using standard operator-intensive methods. New approaches to spacecraft operations are necessary.

---

* Doctoral Candidate, Aeronautics & Astronautics, Stanford University
† Doctoral Candidate, Mechanical Engineering, Stanford University. Caelum Research Corporation, NASA Ames Research Center, CA.

One proposed approach to handling these new constraints is automated health management. This method reduces cost in two ways: operator workload is reduced by performing routine health monitoring using automated systems; and communication requirements are reduced by migrating this automation to the spacecraft. Of course, spacecraft autonomy brings new challenges in determining the amount and type of information necessary to relay to ground operators. For example, operators will no longer be able to access an archived database to "catch up" on spacecraft health history. There are also questions of how robust spacecraft autonomy is to unexpected behavior and events. In meeting such challenges, the role of telemetry in spacecraft operations is being redefined.

After reviewing the laboratory's satellite and operations architecture, this paper outlines the beacon-based health management approach undertaken by the Space Systems Development Laboratory (SSDL) at Stanford University. The ASSET experimental mission operations system is being used to develop methods for automated health management, with the SAPPHIRE microsatellite as a testbed. After background information about the satellite and operations architecture, this paper outlines the health management system. Questions are raised about informing operators of important issues on-board, leading to a new view of telemetry. This paper also highlights the role of universities in researching and developing telemetry systems.

THE SPACE SYSTEM DEVELOPMENT LABORATORY RESEARCH PROGRAM

SSDL was chartered in 1994 to provide world class education and research in all aspects of spacecraft design, technology, and operation. To achieve this goal, SSDL members enroll in a comprehensive academic program composed of coursework, project experience and research investigations. As one of their investigations, SSDL is actively involved in research in spacecraft operations and automation.

The Satellite Quick Research Testbed (SQUIRT) Microsatellite Program - The SSDL SQUIRT program [1] is a yearly project through which students design and fabricate a real spacecraft capable of servicing low mass, low power, state-of-the-art research payloads. By limiting the design scope of these satellites, the project is simple and short enough so that students can see a full project life cycle and are able to technically understand the entire system. Typical design guidelines for these projects include using a highly modular bus weighing approximately 25 pounds, a hexagonal form that is roughly 9 inches high by 16 inches in diameter, amateur radio communications frequencies, and commercial off-the-shelf components. Missions are limited to about one year of on-orbit operation. Since little money is available for operations, a highly automated mission control architecture is being developed.

The Stanford Audiophonic Photographic Infrared Experiment (SAPPHIRE) Microsatellite - SAPPHIRE is the first SQUIRT spacecraft. Its primary mission is to characterize the on-orbit performance of a new generation of infrared horizon detectors, in addition to flying two student payloads (a digital camera and a voice synthesizer). Student research interests are also driving experiments in nontraditional sensing [2] and automated operations. SAPPHIRE is hexagonal, measuring 17" across its longest dimension and 13" high. It is primarily constructed of commercially available equipment, such as amateur radio kits and a Motorola 68000 series microcontroller, with some space-qualified elements, such as batteries. SAPPHIRE is being completed by a core of volunteers and research students, and is currently undergoing environmental testing.

The Automated Space System Experimental Testbed (ASSET) System - The ASSET system [3] is a global space operations network under development within SSDL. The first goal of this system is to enable low-cost and highly accessible mission operations for SQUIRT microsatellites 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 1 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 developmental 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, anomaly management, contact planning, data formatting and distribution, and executive control.

SAPPHIRE and all future SQUIRT satellites will be operated through ASSET. In addition, controllers for a number of other university and amateur satellites have expressed in becoming part of the system. As for ground stations, the Ogden and Stanford ground stations are the first two facilities to be included. Several other stations throughout North America and Europe have been identified for future integration.

BEACON-BASED HEALTH MANAGEMENT DESCRIPTION

A beacon-based health management concept was proposed for a deep-space mission to the planet Pluto [4]. This concept is being prototyped as a part of the SAPPHIRE mission [5]; its main features are summarized, below, and the new elements are further detailed. The signal flow is presented in Figure 2.

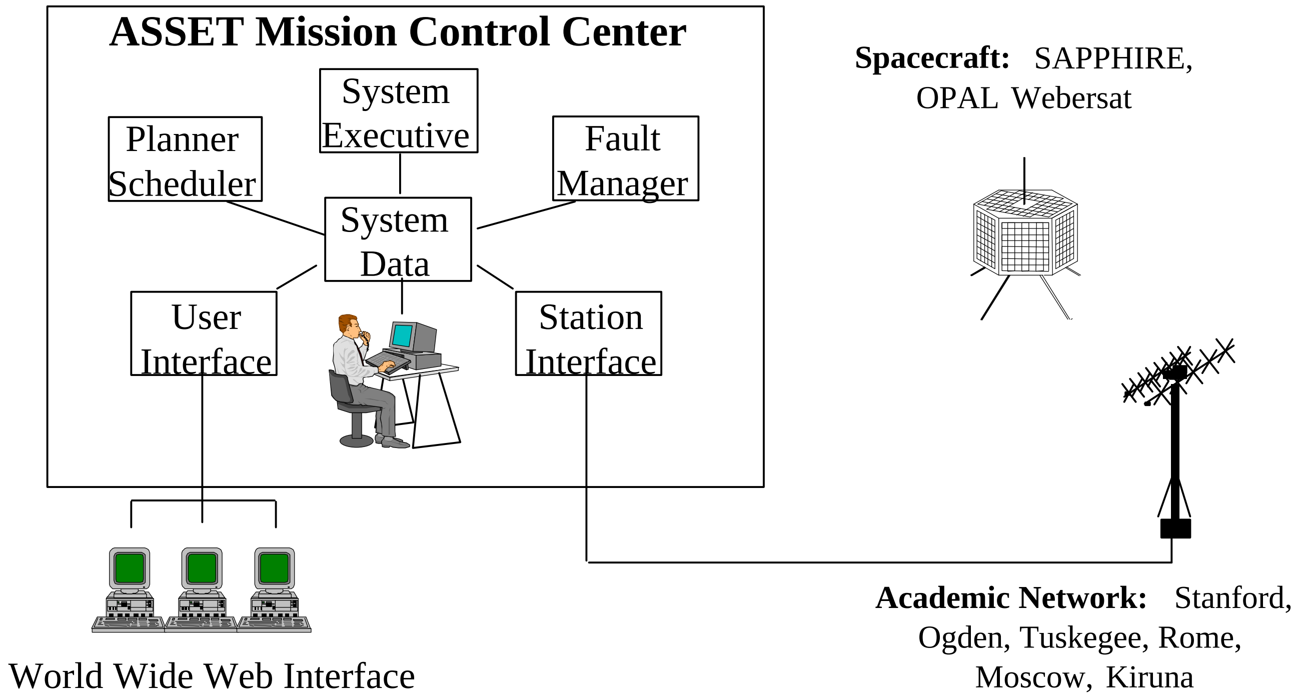

***Figure 1** - The ASSET Space System Architecture*

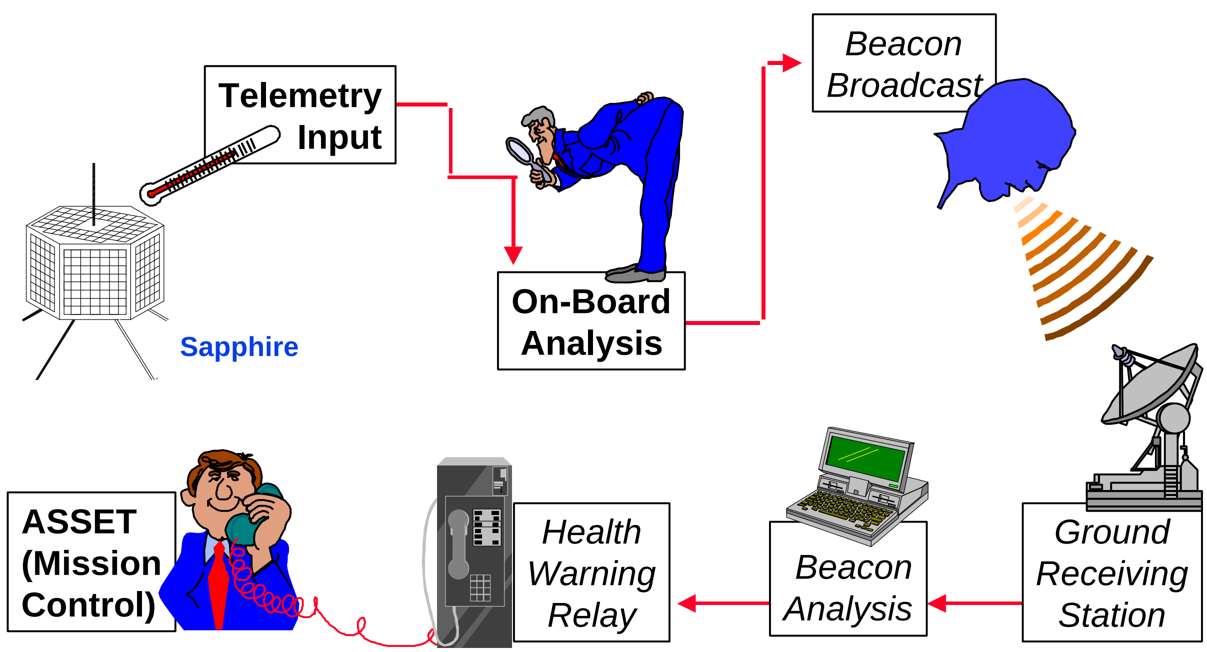

***Figure 2 –** SAPPHIRE Health Beacon Signal Flow*

SAPPHIRE Health Monitoring – SAPPHIRE will monitor its own telemetry sensors, comparing measured values with commandable entries in a state-dependent limit table. Certain measurands will be validated by aggregate analysis. For example, all body-mounted solar panels cannot see the Sun at once; this knowledge is used to help verify that the solar panel current readings "make sense". These modest steps provide SAPPHIRE with an anomaly detection system far more mature than most spacecraft.

Depending on the seriousness of the limit violation, the spacecraft health is assessed to be one of four values. For example, when measurands are within limits, the spacecraft is judged to be Normal. Out-of-limits with moderate impact, such as an overheating transmitter, is considered an Alert. Out-of-limits that can rapidly jeopardize the mission elevates the health status to Critical. Finally, Emergency condition is defined to be an unexplained computer reset. Note that the rules by which measurands trigger the modes, and the limits for each, are defined by the operations team. This is because the four beacon modes are intended to be a mapping from spacecraft state to operator action.

Health Beacon Transmission – The health beacon is transmitted with the AX.25 packet format common to Amateur radio users. SAPPHIRE's main transmitter also serves to broadcast the beacon message, with an output power of around 500mW. The message consists of SAPPHIRE's call sign (part of the protocol) and the health assessment identifier. In this manner, the information from thirty-two analog sensors and six digital lines has been compacted into a few bits that tell an operator whether or not SAPPHIRE can continue to perform its mission. And while such information once had to be collected over time for eventual download and processing at mission control, spacecraft health is now continuously monitored and available anytime the spacecraft is within range of a low-cost receiving station.

Receiving Station – SSDL has partnerships with universities in Alabama, Montana, and Sweden to develop a simple receive-only system for health monitoring. These stations will listen for SAPPHIRE beacon transmissions and notify mission control of the results by electronic mail. It is intended to put these stations at locations around the world, giving SAPPHIRE (and other spacecraft in the ASSET network) near-global coverage for health monitoring.

Mission Control Center – Once mission control receives a beacon monitoring update from a remote station, it logs this information and then takes appropriate action. Depending on the health assessment, there are varied responses, from storing the update in the system database to paging the operator on call and rescheduling the network to contact and recover a failed satellite.

Implementation – Currently, SAPPHIRE's flight software is being modified to include health monitoring and beacon message broadcasts. Beacon receiving stations are being developed by the university partners. Automatic notification from a beacon monitoring update has been demonstrated, and ASSET's scheduling capabilities are still in development. A functional demonstration of the complete system will take place in the fall of 1997.

ENGINEERING DATA SUMMARY

The second challenge in automated health management is to provide the operators with the context necessary to make informed decisions. Normally, these operators spend hours closely watching the various spacecraft measurands, enabling them to spot trends and develop an intuitive feel for the various subsystems. Familiarity with the system is extremely important for such tasks as anomaly management, because not all faults are directly observable by measurands, nor is it always clear which recovery action will least impact other components. However, in this approach to automated health management, the spacecraft is now performing the routine data analysis once handled by operators. In addition, the raw telemetry is no longer available on the ground for offline analysis or operator perusal.

The problem statement for data summary, then, is to reduce the set of on-board sensor data to what is necessary for operators to carry out their duties. This reduced telemetry set is the only information sent back to mission control, apart from science or mission data from the payload. Sending additional information to the ground wastes scarce communication resources; sending less prevents operators from ensuring mission success.

Now, this is not simply an application for data compression, for the goal is not use the summary to reconstruct the complete measurand set. Rather, it is to identify the information valuable to operators and present it succinctly. Numerical compression makes no use of the built-up body of knowledge about a spacecraft. Such knowledge aids the transformation of measurands to into data summaries that are full of meaningful content. Such an approach promises greater data savings over compression and provides operators a head start in such tasks as anomaly management.

The task of anomaly management was chosen as the first candidate for data summary. Once an anomaly is detected, operators must determine what fault there is, if any, and take action to return the spacecraft to as normal a condition as possible. Two candidate solutions have been identified. Both share their roots in a concept of model-based analysis, taking advantage of what is known about the spacecraft components and their interactions in developing summaries. Other, non-model-based approaches exist, such as statistical summaries and curve-fitting, but have not been selected as initial candidates.

The first candidate approach is to log only the out-of-limits measurands. The limits will be defined according to component and system reliability values, and will vary for different operational modes. This is expected to reduce the amount of data sent back by more than a factor of ten[‡]. The shortcoming of this method is that all the data leading up to the anomaly is ignored, so it is difficult to perform trend analysis. Secondly, this method ignores other components that are exhibiting unusual, but not out-of-limits, behavior, which may give hints as to the source of a hidden fault.

The next candidate approach integrates the work done for advanced on-board anomaly detection using artificial intelligence techniques. Such a system not only monitors the on-board sensors, but it analyzes the results to identify faults, using a built-in spacecraft model. In contrast to the first method, the measurand limits are dynamically set models of each component. The advantage to this method is that the same model can be reproduced, which means that the spacecraft needs only to send back the information to update the ground model, such as inputs and the functional status of components. Should an anomaly occur, the spacecraft is responsible for determining the source of the anomaly; this information is considered an update to the spacecraft state and the ground is notified only if this requires a change in operations. It is unclear what the savings in transmitted data will be, though expectations are that this method should out-perform the first candidate while further reducing the responsibilities of the operators. A known drawback is that it requires significant on-board computation, as well as a detailed model of the whole spacecraft. Additionally, it is unclear whether all this computational effort succeeds in paring the data down to that which the operators require. Further study is needed.

THE UNIVERSITY'S ROLE IN TELEMETRY RESEARCH

SSDL's work in automated health management provides insight and examples in how universities can contribute to the research and development of telemetry approaches. Put bluntly, universities have the opportunity to fail – for failure is a fundamental part of the education process. A university laboratory like SSDL can afford to investigate risky projects or to approach problems from widely different angles. This is because the primary outputs of a university are not successful products or methods, but well-trained, capable students. Obviously, understanding why a specific approach did not work is critical to the development of good students. But a failed project is not a true failure if new solutions (and better-informed students) come out of it. Unlike a competitive industrial project, universities can afford to fail.

[‡] An Air Force study [6] determined that the spacecraft was determined to be "normal" in more than 90% of contacts. Admittedly, such decisions involved an operator who could ignore the many false alarms generated by inadequate limit definitions. But this new approach assumes a mode-dependent limit checking, adjustable over time to account for component degradation, which should eliminate most of these false alarms.

FUTURE WORK

Obviously, the highest priority is to complete the development of the beacon-based health monitoring system. The summer will be spent completing and testing SAPPHIRE's flight software, including the on-board health management. Work will continue in coordinating the construction of beacon receiving stations from the detailed designs that have been generated. And ASSET's capabilities will be expanded to allow for operation of SAPPHIRE through its system, culminating in a complete system demonstration in fall 1997.

Once those elements are in place, the value of this automated health monitoring approach can be demonstrated. SAPPHIRE will be operated for a period of a week or more, first using the standard operator-intensive approach, then with the beacon-based monitoring. Three factors will be compared: accuracy of fault detection, speed of response to on-board anomalies, and the cost in operator hours and active satellite-to-ground contacts. Since the automated fault detection methods are identical to the ones used by operators, it is expected that short-term accuracy will remain the same. However, significant improvements are expected in both speed of response and cost, since this passive network will continuously listen over a wide area and operators are required only for anomalous events.

Engineering data summary research is a wide-open area that needs further investigation. The candidate approaches need to be evaluated for performance, and more candidates need to be developed. While data summary will not be a part of SAPPHIRE's flight software, summary algorithms will be implemented on the ground to assess their impact on flight data analysis.

CONCLUSION

The use of a low-power beacon and distributed listening stations is a viable solution to the new challenges in spacecraft health management. Automation of limit detection and migration of this process to the spacecraft will significantly reduce the spacecraft-to-ground communication link and the time spent by operators. This beacon-based health management approach will be demonstrated using the SAPPHIRE spacecraft through the ASSET operations architecture, in cooperation with Stanford's university partners throughout the world.

Given this anomaly detection approach, a primary concern is to assist operators given their inability to access the complete measurand set. Two candidates in the area of anomaly management have been defined: one summarizes all out-of-limits measurements, the other uses advanced fault detection methods to reduce the measurands to an abstraction of the spacecraft operational status and the known inputs. Neither approach has proven to be completely satisfactory, and both approaches will be further refined to develop new candidates.

Work on this research topic highlights the need for better defining the information – both the type and amount – needed for spacecraft operators to perform anomaly management. Answers to that question will also affect the future role of telemetering systems in spacecraft operations. The traditional function of spacecraft telemetering – measuring sensors at a distance – is being challenged by the high cost of remote operations. Data interpretation in the next generation of spacecraft be handled on-board, with only an abstraction of the measurands returned to the ground. Such a profound change in telemetering deserves further study. The university environment, with its greater flexibility and room for failure, is an important contributor to such investigations.

ACKNOWLEDGEMENTS

The authors wish to thank of the SAPPHIRE team for their dedication and efforts, especially the flight software programmers, Rajesh Batra and Kenneth Koller. The Jet Propulsion Laboratory and NASA Ames Research Center are acknowledged for their support and assistance with this research project. Finally, appreciation is extended to all of the institutions and organizations that have participated in, contributed to, and provided feedback concerning various aspects of the ASSET system. This includes all of the previously mentioned schools that are interested in integrating their spacecraft and/or ground stations with the ASSET network. Special thanks goes to the receiving station partners at Tuskegee Montana, and Sweden. This work has been performed in partial satisfaction of graduate studies at Stanford University.

REFERENCES

[1] Kitts, Christopher A., and Robert J. Twiggs, "The Satellite Quick Research Testbed (SQUIRT) Program," Proceedings of the 8th Annual AIAA/USU Conference on Small Satellites, Logan, Utah, September 16-22, 1994.

[2] Swartwout, Michael A., Tanya A. Olsen, and Christopher A. Kitts, “The Omni-Directional Differential Sun Sensor”, Proceedings of the 31st Annual International Telemetering Conference, Volume XXXI, IFT, Las Vegas, NV, 1995, pp. 730-736.

[3] Kitts, Christopher A., "A Global Spacecraft Control Network for Spacecraft Autonomy Research." Proceedings of SpaceOps '96: The Fourth International Symposium on Space Mission Operations and Ground Data Systems, Munich, Germany, September 16-20, 1996.

[4] Wyatt, E. Jay, and John B. Carraway, "Beacon Monitoring Approach to Spacecraft Operations", Reducing the Cost of Spacecraft Ground Systems and Operations, Oxfordshire, United Kingdom, September 27-29, 1995.

[5] Swartwout, Michael A., and Christopher A. Kitts, "A Beacon Monitoring System for Automated Fault Management Operations", Proceedings of the Tenth Annual AIAA/USU Conference on Small Satellites, Logan, Utah, September 16-19, 1996.

[6] Farmer, Mike and Randy Culver, "The Challenges of Low-Cost, Automated Satellite Operations", Proceedings of the 31st Annual International Telemetering Conference, Volume XXXI, IFT, Las Vegas, NV, 1995, pp. 203-209.