

An Overview of the Beacon Monitor Operations Technology

E. Jay Wyatt, Mike Foster, Alan Schlutsmeyer, Rob Sherwood, Miles K. Sue

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California USA
e-mail: e.j.wyatt@jpl.nasa.gov
Tel: 818-354-1414 Fax: 818-393-3654

ABSTRACT

The beacon monitor operations technology is aimed at decreasing the total volume of downlinked engineering telemetry by reducing the frequency of downlink and the volume of data received per pass. Cost savings is achieved by reducing the amount of routine telemetry processing and analysis performed by ground staff. Antenna resources are also conserved so more missions can be supported with existing ground stations. Beacon monitor operations is a process for allowing the spacecraft to transmit a tone instead of telemetry that indicates the urgency of ground intervention. Onboard engineering data summarization and automated ground antenna scheduling are important components of the system design. The technology will be used on upcoming missions to Pluto and Europa and flight validation will occur on the New Millennium Program Deep Space One (DS-1) mission to be launched in July of 1998. The technology can also be adapted for use on earth orbiter missions.

1. INTRODUCTION

The beacon monitor operations technology was conceived in response to a need to lower the cost of mission operations for NASA's planned mission to Pluto. Since the spacecraft was to be highly 'autonomous', it seemed logical to leverage enhanced onboard capability to allow the spacecraft to determine when telemetry downlink should occur. Reducing the frequency and volume of downlinked engineering data would enable cruise operations to be carried out with less staff and would reduce the loading on an already highly constrained Deep Space Network (DSN). [1] [2]

From an operations perspective, it is not appealing to simply mandate infrequent tracking. A better approach is to have the spacecraft continuously transmit one of four tones that indicate how urgent it is to track the spacecraft for telemetry. Since no telemetry is modulated with these tones, the ground process is much simpler and less expensive. Tones provide the necessary assurances that the mission is

proceeding as planned while conditions are nominal. When an anomaly occurs the tone system facilitates quick response, minimizing impact to the mission timeline.

In order to engineer a complete solution, certain ground and flight components should be present in the end-to-end design. One technology component is a capability for generating summaries of engineering data onboard the spacecraft. These summaries are needed to quickly provide operators with data associated with important spacecraft events when telemetry downlink is necessary. Another necessary technology component for NASA deep space missions is a capability for scheduling DSN antennas based on events rather than through pre-negotiated agreements.

The beacon monitor technology has been manifested for flight validation on the DS-1 mission as the Beacon Monitor Operations Experiment (BMOX). This activity is currently the focal point in developing the technology. All of the BMOX components and the operational concept will be validated during DS-1 operations. Validation objectives fall into the following categories: (1) onboard engineering summary data generation and visualization, (2) tone selection, transmission, and detection, (3) multi-mission ground support, and (4) operations concept demonstration and assessment. A detailed set of experiments has been defined. As the mission progresses, BMOX component technologies and the operations concept will be gradually validated and will be made available for use in baseline DS-1 operations.

2. OPERATIONS CONCEPT

Beacon monitor operations occurs when the spacecraft is allowed to transmit tones instead of telemetry. The four tone system does not represent spacecraft state, but rather the spacecraft's assessment of how urgent it is to track for telemetry. For this reason, the basic four tone system described in Table 1 can support most any type of deep space mission.

Tone	Definition
Nominal	Spacecraft is nominal, all functions are performing as expected. No need to downlink engineering telemetry.

Interesting	An interesting and non-urgent event has occurred on the spacecraft. Establish communication with the ground when convenient to obtain data relating to the event. <u>Example</u> : device reset to clear error caused by SEU , other transient events.
Important	The spacecraft needs servicing. Communication with the ground needs to be achieved within a certain time or the spacecraft state could deteriorate and/or critical data could be lost. <u>Examples</u> : solid state memory near full, non-critical hardware failure.
Urgent	Spacecraft emergency. A critical component of the spacecraft has failed. The spacecraft cannot autonomously recover and ground intervention is required immediately. <u>Examples</u> : 1553 bus failure, PDU failure, SRU failure, IPS gimbal stuck.
- No Tone -	Beacon mode is not operating, spacecraft telecom is not Earth-pointed or spacecraft anomaly prohibited tone from being sent.

Table 1 Beacon Tone Definitions

Although the flight project determines how frequently to poll the beacon tone, a rate of once per day is generally considered optimal. When the tone indicates that tracking is required, an antenna is scheduled to retrieve the summary. If the summary reveals that further intervention is necessary, another track is scheduled to uplink commands or retrieve more data. A mission would likely transition gradually to beacon operations once routine operations are underway. The summarization flight software adapts to downlink rate, allowing much larger telemetry summaries in the early phases of the mission. As operators become more comfortable using the summaries, frequency of nominal tracking and summary size can be scaled-back for additional cost savings. Summary passes can be much shorter than the 6-8 hour passes used on traditional deep space missions, even when larger summaries are desired.

3. FLIGHT SYSTEM HARDWARE

The responsibility of flight hardware is to generate and transmit beacon signals representing the four beacon states. In theory, this function can be supported with a very simple transmit subsystem that has a modulator, an exciter, an amplifier, and an antenna. The telecommunications subsystem for space missions in general is more complicated. The complexity varies from one mission to another and is affected by the coverage requirements, the number of frequency bands, and the redundancy requirements. For the DS-1 mission, the telecom subsystem includes a Small Deep Space Transponder (SDST), a 12.5 W X-band solid state power amplifiers (SSPA), a 2.5 W

Ka-band SSPA, two X-band low-gain antennas (LGA), and one dual-frequency X- and Ka-band high-gain antenna (HGA). The gain of the HGA is ~25 dBic gain at X-band and 27 dBic at Ka-band. The gain of the LGA is ~9 dBic. The SDST is a highly integrated package that includes a telemetry modulator, exciter, and command detector. The DS-1 telecommunications subsystem will be used to generate and transmit beacon signals for the experiment, without modifications.

The four beacon messages are each represented by a pair of tones, centered around the carrier. These tones will be generated by phase-modulating the RF carrier by a squarewave subcarrier using 90 degrees modulation angle. The carrier will be completely suppressed. The resulting downlink spectrum will consist of tones at odd multiples of the subcarrier frequency above and below the carrier. The higher harmonics will be ignored; only the tones at the fundamental frequency will be used to represent the transmitted message. Ignoring the higher harmonics results in a slight loss of signal strength. The SDST has the capability to generate a wide range of subcarrier frequency to represent the four beacon messages. However, the downlink frequency uncertainty and detector complexity together constrain the selection of the subcarrier frequencies. The frequency uncertainty is caused by a combination of on-board temperature variations and uncorrected residual Doppler frequency. For the Beacon Monitor Experiment, the four subcarrier frequencies are 20, 25, 30, and 35 kHz. The signal structure is shown in Figure 1.




Figure 1 Signal Structure

4. FLIGHT SYSTEM SOFTWARE

The amount by which beacon monitoring reduces operations cost depends largely on the level of autonomy achieved onboard the spacecraft. Systems that can perform more robust recovery from anomaly conditions and provide flexible onboard data management can achieve the most benefit from beacon operations. The two primary flight software innovations implemented through the beacon monitor development effort are onboard engineering data summarization and beacon tone selection. The tone selection module is a software component that implements the functionality required to select tone states based on spacecraft health information. The summarization module is a comprehensive

architecture for creating summaries of engineering data between telemetry downlink periods.

4.1 Tone Selection

The tone selector module maps fault protection messages to beacon tone states. This module outputs the tone state as a telemetry packet and a message to the Small Deep Space Transponder to actually set the tone. The output of the module can be turned OFF, ON or RESET by issuing a ground command. Tone state always transitions to a higher level of urgency until reset by a ground command. The tone selection module for DS-1 currently consists of 1583 source lines in C and uses 31 Kbytes of memory when compiled for a Power PC series processor. The processing load is all in response to external messages which come at undetermined intervals (less than 1/sec). The processing time for one message is negligible.

4.2 Data Summarization

Flight software for summarization integrates several techniques into one cohesive architecture for providing operators with information required to analyze spacecraft engineering data. The method for creating the summary data is highly event-driven and uses the data prioritization capabilities built into the telemetry management software. It has been a design goal to integrate the most advanced techniques possible into the architecture. Transforms and adaptive alarm thresholds are key components of an architecture creating top-level summary statistics, episode data (high-resolution culprit and causally related data), low-resolution "snapshot" telemetry, and user-defined data .

Transforms

Each sensor will use a selected subset of the five transforms included in the summarization software package. Those transforms are: minimum value, maximum value, mean value, first derivative, and second derivative. To compute the transforms, the raw data from each sensor is stored in a lag vector of a predetermined length, and the value of the transform is computed across the vector. It should be noted that the first and second derivative transforms are not true derivative functions. The first derivative transform simply computes the rate of change between the first and last points in the vector. Similarly the second derivative transform calculates the rate of change between the difference of the first two values in the vector and the difference of the last two values in the vector. Minimum, maximum, and mean values are also calculated for the first and second derivative transforms. Transforms require more computational cycles than straight limit comparison; however, the

time required to perform the mathematical operations is inconsequential.

Using the transforms to determine when an episode should begin provides more flexibility than simply testing the raw data against red-line limits. If an episode were to begin whenever a sensor value moved above/below a certain value, then red-line limit sensing would be sufficient. This functionality, which is adequate for many sensors, is still provided when we use the transforms. By using the maximum and minimum value transforms, we can duplicate the red-line limit sensing. If the maximum/minimum value goes above/below the set boundary, then an episode is started. However, using the transforms also provides this functionality while better handling some cases, where the traditional method would either fail or signal false alarms.

In traditional limit sensing, there is no concept of memory or time in examining values, i.e. an episode cannot be triggered if a value stays above/below a certain level for a certain amount of time. However, this type of behavior can be captured with the mean value transform. A value falling below a limit for too long will reduce the mean, causing an episode to begin. With traditional methods, an episode would occur anytime the value dropped below the threshold, possibly leading to many false alarms.

During nominal spacecraft operation, some sensors may change value in a cyclical manner. Traditional limit sensing would be able to detect when the sensor went above or below its expected values, but would fail to detect the anomalous situation where the value did not change. To handle this case, we could use the first and second derivative transforms. For a constant value, both first and second derivative functions would be zero, signaling that an episode should be triggered.

Adaptive Alarm Thresholds [3]

The current state-of-the-art in anomaly detection is to use limit-sensing, in which the current sensor value is compared against predetermined high and low "red-lines". Such red-lines are typically constants across many or all mission modes and it is difficult to determine tight limits which will work well throughout the mission. Thus, to avoid frequent false alarms, the red-lines are made imprecise, leading to missed alarms and missed opportunities for early anomaly detection.

An alternative to red-line limits is envelope functions learned from historical and/or simulated data. Limits become dynamically changing values instead of static constants. These limits are functional values based upon the values of related sensors and other factors, such as the current operational mode of the spacecraft. Although learning precise envelopes can take longer than determining red-lines, initial loose envelopes can be learned quickly. With further training, the bounds can be incrementally tightened, while still retaining a low false alarm rate. Since the

learned envelopes are tighter than red-lines, they have a much lower missed alarm rate. Novel training methods are being employed to avoid bounds which cause alarms in nominal training data. Therefore, these envelopes are loose enough to avoid false alarms, provided the training and validation data are representative.

In order to learn the envelopes, the ELMER (Envelope Learning and Monitoring via Error Relaxation) algorithm will be used. For training purposes, ELMER can be run on the ground using historical spacecraft data, examining both anomaly and nominal data sets in order to determine accurate bounds. For certain phases of the DS-1 beacon monitor experiment, the ground trainer will produce limit functions for uplink as shown in Figure 2.




Figure 2 Ground Trainer Software

Flight Software Architecture

The summarizer consists of 3 subroutines: data collection and processing, episode, and mission activity. The data collection and processing module receives data from various domain units via C function calls and applies summary techniques to these data, producing summary measures for downlink to the ground. The Sampler/Summarizer/Episode Identifier module is awakened once per processing interval (1 sec. for DS-1), at which time it reads the current value of all the spacecraft data items available, computes derived data such as running minimum, mean and trends, checks the raw and derived data against upper and lower limits, starts and ends "episodes" describing out-of-limits conditions, and produces historical and episodic telemetry. It also responds to spacecraft state messages to determine the current spacecraft activity, which determines which set of bounding limits to use. The architecture diagram shown in Figure 3 shows the interfaces between the subroutines and the major constituents of the summaries.

Data collection is performed by providing a central memory array accessible to all flight software. The various flight software modules update the memory array once per second or whenever the data changes, whichever is less often. The update is done by direct write, thus the time for a single update is just a few microseconds.

The episode subroutine looks for anomalies within the data and summarizes all data relevant to the anomaly. The episode subroutine receives summary

and engineering data internally from the summarizer/sampler module and compares the data with alarm limits. If the limit is exceeded, the subroutine spawns a new episode and collects relevant data from the summarizer/sampler. This data is in the form of one minute summaries that start five minutes before the episode and end five minutes after the end of the episode. At the end of the episode, the subroutine outputs episode name (out of limit data ID), high limit, low limit, relevant data, start and end times.

The mission activity subroutine determines the overall spacecraft mode of operation. This determination is used to select data and limits for a particular episode in the episode subroutine. For example, solar electric propulsion (SEP) sensor values may be important while using SEP, but if the satellite is in RCS control mode then SEP sensor values could be ignored. In addition, the ACS rate limits might be different during cruise than during a maneuver. As this example points out, it is necessary to use the mission activities to determine which data to use for episode identification. The mission activity is intended to be exclusive. When a new mission mode start message is received, the previous mission mode is assumed to have ended.

The sampler module and its related data gathering module currently consist of 3038 lines of source code and 222 Kbytes of memory on the Power PC series processors. Activity determination is a rare event and processing time is negligible. The once-per-wake up processing time for DS-1 averages out to 30ms.

5. TONE DETECTION SYSTEM

The ground monitor station is fully automated and its operation is driven solely by schedule and predicts. The received signal is first down-converted to IF, sampled, digitized, and recorded. The digitized signal is processed by the signal detector, which performs a non-coherent detection using the Fast Fourier Transform (FFT). The detected signal is then decoded by the Message Decoder, and the decoded message is then disseminated to the mission operations team and other users. A block diagram of the station is shown in Figure 4.




Figure 4 Monitor Station Block Diagram

The signal detector contains four tone detectors, one for each message. To insure proper signal detection, the bandwidth of each tone detector must be

sufficiently large to accommodate the frequency uncertainty and frequency drift of the downlink frequency, i.e., the beacon tones for a given message will not drift outside of the passband of the detector for that message. The FFT is employed to compute the energy of all spectral pairs having spacing corresponding to the four beacon signals. Because of oscillator instability, Fourier transforms cannot be produced over long time intervals. The total observation time is divided into short intervals. FFTs are first performed over these short intervals and then incoherently combined after the frequency drift has been removed. The maximum of the outputs of the four tone detectors is then selected and compared against a pre-determined threshold to determine which message has been received. A block diagram for the signal detector and the message decoder is shown in Figure 5.

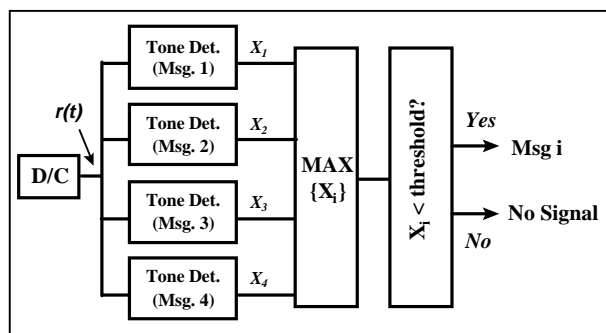


Figure 5 Monitoring Signal Detector and Message Decoder

6. TONE HANDLING & REPORTING

The beacon signal detection and message delivery system for the DS-1 experiment is shown in Figure 6. DSS-26 (a 34m DSN antenna) will double as a monitoring station as well as a demand access station. The beacon message is first received and decoded by the remote monitoring station and subsequently transmitted to the BMOX team at JPL via a secured link, such as the NASA Science Internet. BMOX in turn forwards the beacon message to DS-1 Mission Operations and other end users, including the Demand Access Scheduler, using e:mail or pagers. Depending on what message has been received, different activities will be carried out by the BMOX team, the Demand Access Scheduler, the mission operations team, and the DSN station. If the tone indicates that telemetry tracking is required, the Demand Access Scheduler will schedule a downlink track for the demand station to receive telemetry from the spacecraft. The scheduler will notify the BMOX team of the schedule. BMOX will in turn notify the Mission Operations team and obtain its approval to carry out the downlink track triggered by the beacon message. One round-trip-light-time prior to the downlink track, a command will be transmitted to the spacecraft by the demand access station or by another 34m antenna station to

initiate the downlink pass. The downlink telemetry will be received by the demand access station, forwarded to the Mission Operations and the BMOX teams, and analyzed.

7. SUMMARY DATA VISUALIZATION

Moving to a paradigm where downlink is infrequent requires new approaches for data visualization on the ground. The onboard summarization architecture provides data at variable resolution based on total available bandwidth and the number of significant episodes since the last downlink. The operator will need to quickly locate the high resolution episode data and would likely use the low resolution (snapshot) data for gaining overall system context. An incremental development process is being used with an end vision to develop automated software that searches the data for important information identified in the downlink and guides the operator through analysis of that data.

8. ANTENNA SCHEDULING [4]

Current missions using the DSN negotiate tracking schedules well in advance of the launch date. While this approach is adequate for missions with pre-defined tracking requirements, it does not mesh with the demand-access paradigm of the beacon monitor approach. Since beacon monitoring requires that the spacecraft initiate tracking, antenna schedules must be formed adaptively and must accommodate varying degrees of urgency in beacon tone states. The DSN advanced technology program is supporting beacon monitor operations development by implementing a method for automated scheduling and providing dedicated antenna resources that can be used for tone detection, telemetry acquisition and command uplink. Segregating pre-scheduled missions from beacon monitor missions avoids conflicts in scheduling paradigms while at the same time evolving a long-term capability within the DSN to support beacon monitor operations.

CONCLUSION

In today's environment, concentrating on reducing downlink is a viable, low risk approach to low cost operations. The technology represents an important, yet practical step in creating advanced mission system designs to achieve cost reduction goals. An added benefit is that the approach is highly compatible with ongoing work in other areas of spacecraft autonomy and the benefits from using beacon monitor operations are enhanced as space systems become more robust.

ACKNOWLEDGMENTS

The research 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] Wyatt, E.J. and J.B. Carraway (1995) "Beacon Monitoring Approach to Spacecraft Mission Operations" 1st International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations, Oxfordshire, UK.

[2] Staehle, R.L., et.al (1997) "Pluto Express: Advanced Technologies Enable Lower Cost Missions to the Outer Solar System and Beyond"

[3] DeCoste, D. (1997) "Automated Learning and Monitoring of Limit Functions", i-SAIRAS '97.

[4] Chien, S., R. Lam, and Q. Vu, (1997) "Resource Scheduling for a Network of Communications Antennas", Proceedings of the IEEE Aerospace Conference, Aspen, CO.

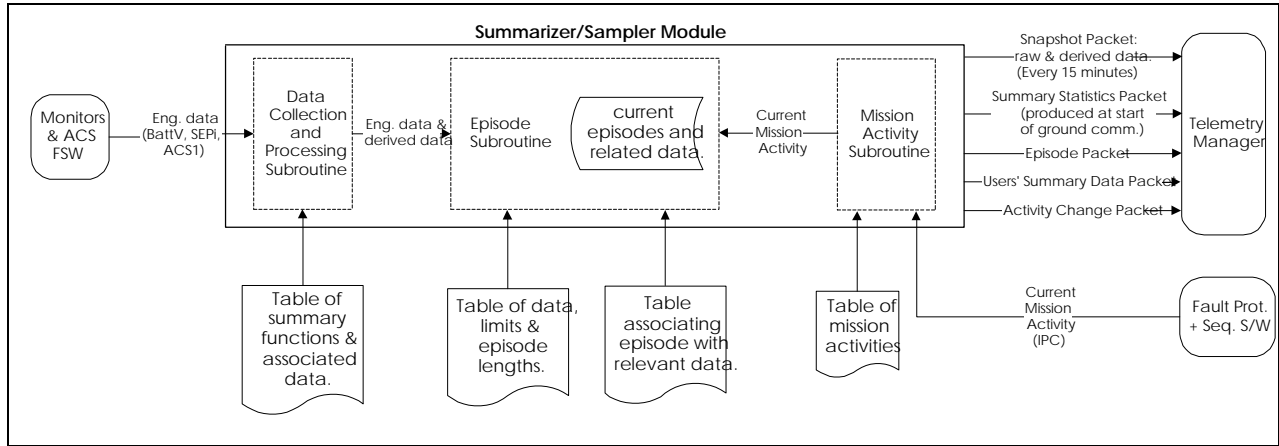


Figure 3 Onboard Engineering Data Summarization Architecture

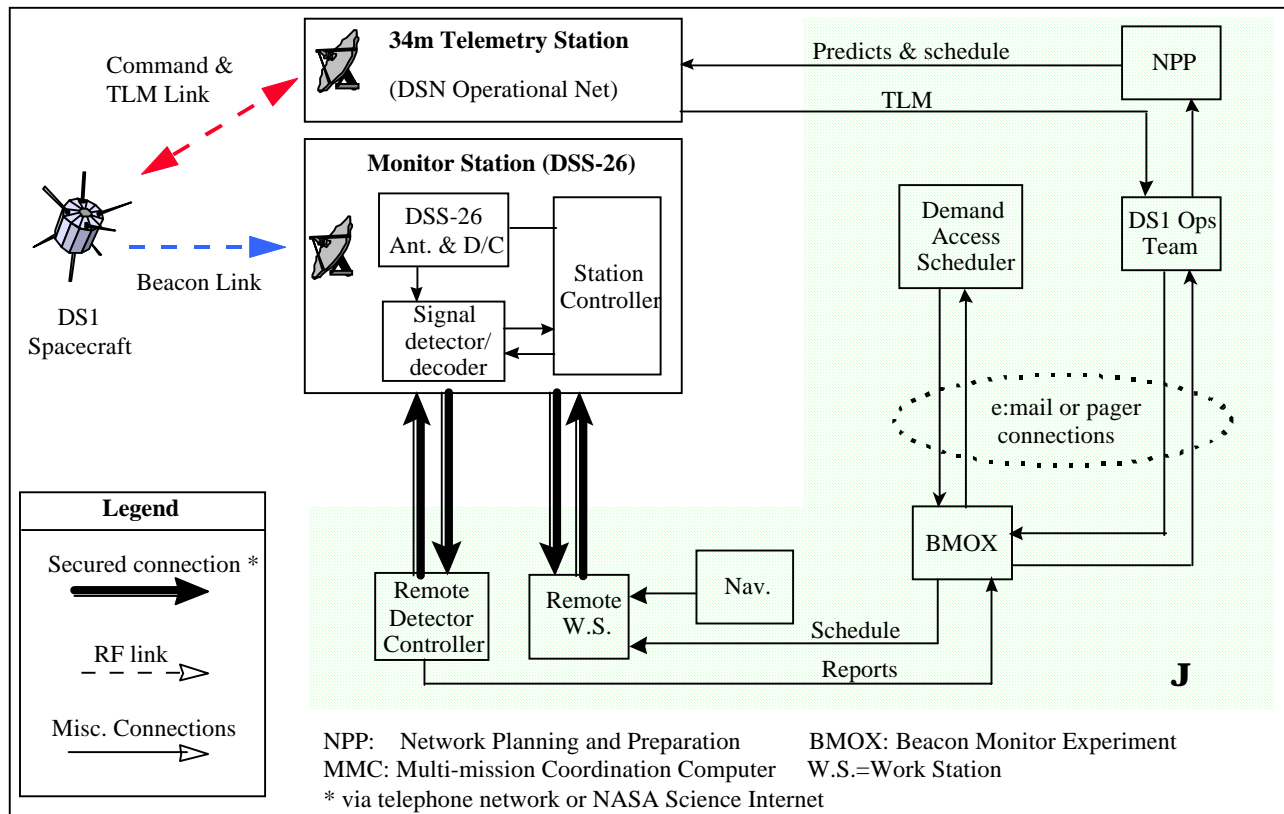


Figure 6 Signal Detection and Message Delivery System