

Lessons Learned During Implementation and Early Operations of the DS1 Beacon Monitor Experiment

ROB SHERWOOD, JAY WYATT, HENRY HOTZ,
ALAN SCHLUTSMEYER, MILES SUE

Jet Propulsion Laboratory, Pasadena, CA, U.S.A.
California Institute of Technology

Abstract. A new approach to mission operations will be flight validated on NASA's New Millennium Program Deep Space One (DS1) mission which launched in October 1998. 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 are achieved by reducing the amount of routine telemetry processing and analysis performed by ground staff. The technology is required for upcoming NASA missions to Pluto, Europa, and possibly some other missions. With beacon monitoring, the spacecraft will assess its own health and will transmit one of four beacon messages each representing a unique frequency tone to inform the ground how urgent it is to track the spacecraft for telemetry. If all conditions are nominal, the tone provides periodic assurance to ground personnel that the mission is proceeding as planned without having to receive and analyze downlinked telemetry. If there is a problem, the tone will indicate that tracking is required and the resulting telemetry will contain a concise summary of what has occurred since the last telemetry pass. The primary components of the technology are a tone monitoring technology, AI-based software for onboard engineering data summarization, and a ground response system. In addition, there is a ground visualization system for telemetry summaries. This paper includes a description of the Beacon monitor concept, the trade-offs with adapting that concept as a technology experiment, the current state of the resulting implementation on DS1, and our lessons learned during the initial checkout phase of the mission. Applicability to future missions is also included.

Keywords: beacon operations; low-cost operations; data summarization; tone messaging; lessons learned.

Abbreviations: DSN Deep Space Network; DS1 Deep Space One; NASA National Aeronautics and Space Administration; AI Artificial Intelligence; BMOX Beacon Monitor Operations Experiment; IPS Ion Propulsion System; SIRTf Spacecraft Infra Red Telescope Facility; OPNAV Optical Navigation.

1. Introduction

1.1 New Millennium Program and DS1

The New Millennium Program, with its advanced technology focus, is one of NASA's many efforts to develop and test an arsenal of cutting-edge technologies and concepts. Once flight proven to work, these technologies will be used by future missions to probe the universe. Deep Space 1, which launched on October 24, 1998, is the first in a series of deep space and Earth-orbiting missions that the New Millennium Program will conduct to demonstrate new technologies in a space-borne testbed. DS1 contains 12 new technologies including an ion-propulsion system (IPS), concentrated solar array, small deep space transponder, and beacon monitor operations experiment (BMOX). It was decided early in the New Millennium Program's conception of missions that a complete testing and proving out of the technologies would require flying them on missions that bore strong resemblance to science missions of the future. Another constraint on the missions derives from the need to respond promptly to future users about whether or not the technologies work in space. It was decided with DS1 that the primary mission should last no longer than about one year. This would allow

sufficient time to conduct an exciting mission and to exercise the technologies under a wide range of conditions without forcing eager potential users to wait unreasonably long before being confident about their use.

DS1 is scheduled to fly by the near-Earth asteroid 1992 KD on July 28, 1999. The primary mission ends on September 18, 1999, by which time DS1 will have completed its mission of demonstrating new technologies. At that time, it may be placed on a new trajectory to encounter comets Wilson-Harrington and Borrelly.

1.2 Beacon Monitor Technology

The beacon operations concept was conceived about four years ago as a method of reducing the cost of operating a deep space mission. The traditional method of operating deep space missions is to contact the spacecraft frequently and download both real-time and recorded performance data. This data was generally acquired daily or several times a week. The problem with operating a spacecraft this way is its cost, both in terms of spacecraft personnel and ground communication resources. Spacecraft engineers spend lots of time analyzing the received spacecraft data. The majority of the time the data shows nominal operation of the spacecraft. The ground station typically used to receive deep space mission data is the Deep Space Network (DSN). There are a limited number of antennas available in the DSN. Using existing technologies and practices, these antennas do not have the capability to support the large number of future missions planned over the next few years. For many deep-space missions, the low bandwidth of the beacon signal and low threshold receiver allow reception using a smaller aperture antenna. These antennas are typically used for Earth-orbiting spacecraft, and there are many available.

The beacon operational concept involves sending one of four simple tones that represent the urgency of contacting the spacecraft for telemetry data. The analysis of spacecraft performance data is accomplished on-board the spacecraft using the beacon and fault protection software. The tone representations are defined by each mission, but generally range from "everything is O.K." to "emergency, contact immediately." Perhaps the best way to think of the beacon message is that of the spacecraft sending a request to the ground that informs the ground personnel how urgent it is to track the spacecraft for telemetry. Thinking of beacon monitoring in this way forces a paradigm shift over the way we traditionally approach operations. Our approach is one where telemetry is only transmitted when it is necessary for ground personnel to assist the spacecraft or otherwise very infrequently if the spacecraft is able to go long periods (a month or so) without requiring ground assistance. The beacon operational concept can be applied to earth orbiting spacecraft and can also be used to facilitate return of science data on missions with adaptive onboard science data processing.

Another piece of the beacon concept is on-board data summarization. If a problem occurs on the spacecraft, ground personnel will need to use past data for analysis. It is impractical to store all performance data on the spacecraft since the previous downlink. In addition, it would take too long to send all this data to the ground after a problem has occurred. The beacon software generates intelligent data summaries. When telemetry tracking is necessary, the intelligent data summaries contain the most relevant information and a complete picture of spacecraft activities since the last contact. The key challenge here has been developing an architecture that enables the spacecraft to adaptively create summary information to make best use of the available bandwidth as the mission progresses such that all pertinent data is received in one telemetry pass.

The primary objectives of this technology are to lower total mission cost and to decrease the loading on DSN antennas. The fact that NASA full-cost accounting requires that new missions pay for tracking cost is a major motivating factor for finding innovative approaches to operations. The following are major themes in the operational concept:

- Substantially reduce the frequency of telemetry tracking during routine operations
- Enable the spacecraft to determine the frequency of contact
- Accommodate varying levels of onboard autonomy (beacon monitoring works for missions with high levels of autonomy as well as for traditional mission designs)
- Conduct operations using shared or on-demand operations teams
- Decrease the size of operations teams

The DS1 spacecraft was chosen to validate the beacon technology for use on future spacecraft missions. Although this may result in lower costs for the DS1 operations team, the primary goal is ensuring that an end-to-end beacon system is available for future deep space missions.

2. DS1 Tone Monitoring Technology

The tone monitoring technology consists of generation, transmission, and detection of the tone signals. There are four tone signals; each uniquely represents one of the four urgency-based beacon messages. The DS1 tone definitions are summarized in *Table 1*.

Table 1. DS1 Tone Definitions

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. <u>Examples</u> : device reset to clear error caused by SEU, other transient events.
Important	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> : 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> : 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.

Urgent Beacon tones on DS1 are sent when the spacecraft fault protection puts the spacecraft in standby mode. This condition occurs when the fault protection encounters a fault that it cannot correct. Standby mode halts the current command sequence, including IPS thrusting.

During the DS1 tone experiment, the Beacon tone is sent regularly at a prescheduled time, i.e., 30 to 60 minutes per day. The Beacon tone is not operated continuously because DS1 requires as much power as possible for IPS thrusting and the tone transmission uses some of the thrusting power.

The signal structure is shown in *Figure 1*. A pair of tones centered about the carrier represents each message. These tones are generated by phase-modulating the RF carrier by a squarewave subcarrier using 90 degrees modulation angle. The carrier (f_c) is completely suppressed. The resulting downlink spectrum consists of tones at odd multiples of the subcarrier frequency above and below the carrier. For the DS1 experiment, the four subcarrier frequencies (f_1, f_2, f_3 , and f_4) are 20, 25, 30, and 35 kHz. Different frequency allocations can

be assigned to different missions. The monitoring system is designed to achieve a low detection threshold. The goal is to reliably detect the monitoring messages with 0 dB-Hz total-received-signal-to-noise-spectral-density ratio (Pt/No) using 1000 seconds observation time.

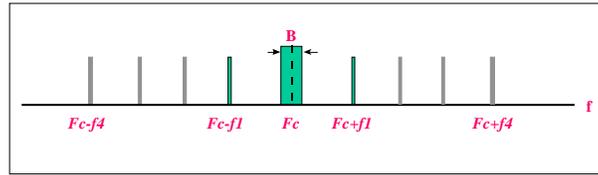


Figure 1. Signal Structure

B=Frequency uncertainty f_c =Carrier frequency
 f_i =Subcarrier frequency for the i^{th} message

The beacon message is first received and decoded by the Goldstone site and subsequently transmitted to the beacon monitoring team at JPL. Next, the beacon message is forwarded to DS1 Mission Operations and other end users, including the Demand Access Scheduler, using email or pagers.

3. DS1 Data Summarization Technology

If the beacon tone indicates that tracking is required, the onboard summarization system provides concise summaries of all pertinent spacecraft data since the previous contact. The summarization system performs three functions: data collection and processing, mission activity determination, and episode identification. The data collection subroutine receives engineering data from the engineering telemetry system via a function call and applies summary techniques to these data, producing summary measures for downlink to the ground. The mission activity subroutine determines the overall spacecraft mode of operation. This determination is used to choose the appropriate data and limits monitored by the episode subroutine. The mission activity is intended to be exclusive. When a new mission activity starts, the previous mission activity is assumed to have ended. The episode subroutine combines summary and engineering data received internally from the data collection subroutine with the mission activity received from the activity subroutine and compares the data with mission activity specific alarm limits. For example, Ion Propulsion System (IPS) sensor values may be important while using IPS, but if the spacecraft is in Reaction Control System (RCS) control mode then IPS sensor values could be ignored. In addition, the attitude rate limits might be different during cruise than during a maneuver. As these examples point out, it is necessary to use the mission activities to determine which data to use for episode identification and to identify the limits of these data. If the limit is exceeded, the subroutine spawns a new episode and collects past relevant data from the data collection subroutine. The past data collected will be one-minute summaries that go back in time as far as the user has defined. (So a five-minute episode would contain summaries starting five minutes before the episode to five minutes after the episode.) At the end of the episode, the subroutine outputs data to the telemetry subsystem for downlink.

Three different types of summarized data are used: overall performance summary, user-defined performance summary, and anomaly summary. Six different telemetry packets have been defined to contain this information. (See Table 2.) The performance summaries are generated at regular intervals and stored in memory until the next telemetry ground contact. They are computed by applying standard functions, such as minimum, maximum, mean, first derivative, and second derivative, to the data. The summarized data are chosen so the spacecraft state can quickly be determined. User-defined summary data are used for obtaining detailed insight into a particular subsystem and are output at the user's discretion. Anomaly

summary data (episodes) are created when the raw and summarized data violate high or low limits. These limits are determined by the subsystem specialist and stored in a table on-board the spacecraft. The limit tables are based on the current mission activity.

Table 2. Summarization Telemetry Packets

Telemetry Name	Description	Output Frequency
Activity	Current value of mission activity	Output on change
Data Sample	Records a snapshot of every raw and summarized data channel	Regular interval, i.e., 15 min.
Episode Summary	Records general data about an out-of-limits data condition called an "episode"	One per episode
Episode Channel	Records specific data about a single data channel's behavior during an episode	One or more per episode
Tone Change	Current state of the beacon tone	Output on tone change
Channel Summary	Summary data about a single data channel's behavior since the last downlink	One for each channel out of limits
User Summary	A user-specified packet containing raw and/or summarized data	Duration user-specified

The software also has the capability to use AI-based envelope functions instead of traditional alarm limits. This new form of event detection will be evaluated in addition to using the project-specified traditional alarm limits. DS1 spacecraft fault protection will only be based on project-specified static alarm limits but the summary data can be generated based on the adaptive limits. Envelope functions are essentially adaptive alarm limits learned by training a neural network with nominal engineering data. The neural net can be onboard or on the ground. For DS1, envelope functions are trained on the ground and then uploaded to the spacecraft.

Tone state and engineering data summaries are displayed on the ground using a special graphical user interface (GUI). The GUI includes a timeline showing all tone changes (detected and telemetry), mission activity changes, data sample packets, downlink summaries, episode data, and user summary data. This type of display environment provides a new approach to interacting with telemetry. The basic idea is that the operator should be able to quickly locate important information in the downlink file. If the onboard summarization system is functioning correctly, the most important information will be available at a high enough sample rate that the operator can perform diagnosis.

4. Tone Response System

The ground response system processes beacon tone messages, notifying appropriate personnel quickly to facilitate interaction with the spacecraft. The system developed for demonstration on DS1 is an early prototype that serves the immediate needs on DS1 and also addresses many of the issues associated with developing a system that can serve multiple flight projects. In general, when a beacon track occurs the track will be logged and someone will be notified. The form of the notification and its latency depends on the perceived urgency of the event. Email will be used for routine events, pager used for significant events requiring prompt attention, and perhaps a synthesized voice call being used for emergencies. Depending on the degree of trust the project has in the notification mechanism it may automatically request antenna time for regular telemetry or emergency tracking. To determine the kind of notification required, events are filtered by urgency and type. All notifications must be

acknowledged, and the time allowed for the acknowledgment should be configurable on a per-project basis.

The project's interpretation of the signal importance will depend on its operations goals. There are two possible interpretations here. First, the mapping of spacecraft-state to urgency of response may evolve as the mission progresses. Early in the prime mission, for example, a device reset may be considered "urgent" because it is wholly unexpected or the consequences are not completely understood. That same event later in the mission, however, may not be considered as urgent and may only trigger the "important" or "interesting" tone. These mappings of spacecraft state to urgency of response can be changed easily by reconfiguring the lookup table. The other interpretation has to do with how each mission defines the latency of response for each tone message. These would vary from mission to mission and may also evolve within a single mission as the operational goals change.

5. Lessons Learned

5.1 Ion Propulsion Missions

The utilization of the ion propulsion system (also called solar-electric propulsion) on DS1 offers an additional advantage in using beacon monitoring. The IPS provides continuous thrust for much of the cruise phase. The operational margin for IPS thrusting represents the duration for which IPS could be off and still allow the spacecraft to reach the target asteroid. Due to the low thrust associated with IPS and because actual thrusting did not start until several weeks after launch, the operational margin is only a few weeks. Telemetry downlink passes are becoming less frequent as the DS1 mission progresses. Eventually, there will only be one telemetry pass per week. If the spacecraft experiences a problem that requires the standby mode, the IPS engine will be shut down. It could be up to one week before the flight team has visibility to that standby mode. Using the beacon tone system during the periods between scheduled telemetry downlinks can be a cost effective way to decrease mission risk because it reduces the likelihood of losing thrusting time and not making the intended target. Other future IPS missions have taken note of this fact and requested beacon tone services to lower their mission risk.

5.2 Software Testing

It was decided to redesign the DS1 flight software about 18 months before launch. This decision greatly compacted an already full schedule to complete the software. As a result, the testing of all non-essential software functions was delayed until after launch. The beacon experiment was considered a non-essential piece of software and therefore was only tested pre-launch for non-interference with the other flight software. In post launch testing, a few problems were discovered that prevent us from starting the beacon software until a new version can be uploaded in February 1999. This is the age-old lesson learned of performing system testing on the software prior to use. Unfortunately, the DS1 schedule would not allow us to do this until post launch.

5.3 Fault Protection Integration

Before the software redesign, the beacon software was tightly integrated with the DS1 fault protection software. The decision was made after the redesign to de-couple the two pieces of software. Previously, the beacon tones were triggered by the fault protection monitors. After the redesign, the mapping of faults to tones was performed using two different methods. All spacecraft standby modes are now mapped to the urgent beacon tone. The interesting and important beacon tones are mapped using beacon software determined limits. Decoupling the fault protection software from the beacon software gives us maximum flexibility to determine what sensors to monitor. Unfortunately, our algorithms for determining faults are not nearly as sophisticated as the fault protection monitors. These monitors can look at many different

values based on conditional logic before determining what fault has occurred. Future spacecraft designed to use beacon operations should plan on completely integrating the beacon tone software with the fault protection software.

5.4 Beacon Signal Frequency Stability

The signals used for beacon monitor are characterized by three things: (1) the signal strength can be extremely low, (2) the initial tone frequencies, which are derived from an on-board auxiliary oscillator, are not known exactly, and (3) the tone frequencies are constantly drifting. The tone detector is designed to detect these types of signals with a high-level of confidence. The maximum frequency uncertainty and the maximum frequency drift rate for the tone detector were established using a Galileo spare transponder. An operational issue was encountered with the DS1 beacon experiment: how and to what extent can we stabilize the temperature of the auxiliary oscillator before the start of a beacon pass? Stabilizing the temperature will reduce the frequency uncertainty and frequency drift, making it easier for the tone detector to detect the beacon signal. Based on data provided by the DS1 telecom personnel, the auxiliary oscillator temperature can undergo a wide range of changes after an OPNAV (optical navigation) maneuver. This results in a very large frequency uncertainty and a very high rate of change (>6 Hz/sec), both of which would exceed the limits of the tone detector (when the signal level is low).

One solution to overcome the OPNAV-related problem is to wait for the transponder temperature to stabilize. Studies by the DS1 telecom personnel indicated that about four hours are needed for the transponder temperature to stabilize after running the OPNAV activity. This operational constraint would not have much impact on the spacecraft and is believed to be the simplest, lowest-cost solution to this problem. We recommend this procedure to improve weak-signal detection for DS1 and future missions using Beacon Monitor.

5.5 Beacon Operations Paradigm

The beacon software makes determinations of spacecraft anomalies. The data summarization component of beacon attempts to summarize related data from these anomalies. These determinations are based upon high and low limits on sensor data. It is important to involve the spacecraft subsystem engineers in the determination of which data to monitor and the setting of the limits on these data. They are the personnel most familiar with the operational characteristics of each subsystem and therefore should be determining interesting and fault conditions for their subsystem. Also, by involving them in the data summarization definition, they will become better acquainted with the beacon software and will be more inclined to use it during crisis situations.

5.6 Other Possible Implementations

Earlier it was stated that the lack of a beacon tone implied there was a problem with the telecommunication system or beacon software. It's also possible to consider non-detection a good response since an autonomous spacecraft may be doing something more important than just telling the ground it's OK, but that is not true indefinitely. If you don't detect the spacecraft for some number of days then you have a problem. In other words, time since previous tone and tone history are both necessary to interpret the beacon tone.

There is another proposed beacon concept for an earth trailing spacecraft (SIRTF) that involves using one tone. SIRTF plans to track every 12 hours, but would like to have beacon tracking every 2 hours. The idea is that the spacecraft would only send a beacon tone if it had a problem. The possible beacon detections are 1) help tone, or 2) no detection. Normally the spacecraft would be busy doing observations, but if it had a problem it would turn to earth point and start transmitting a carrier signal. This beacon signal could shorten the anomaly

response time from 12 hours to a maximum of 2 hours. This requires no modification to the already designed spacecraft since there is no need to distinguish fine levels of urgency. SIRTf management considers this important because their design does not include a transponder that supports beacon tones. There is one drawback with this operation. When the tone detector fails to detect a beacon signal, one can not tell whether (1) the spacecraft is fine and no beacon has been transmitted, or (2) the spacecraft has an anomaly and fails to transmit.

6. SUMMARY

Beacon tone operations can be used to lower the cost of operating space missions while simultaneously decreasing their risk. The concept involves a paradigm shift from routine telemetry downlink and ground analysis to onboard health determination and autonomous data summarization. This technology is being tested on DS1 and is required for several future deep space missions. IPS missions gain an added advantage of power savings from reduced telemetry downlinks and the associated increased thrusting time. Beacon operations will enable more of the smaller, more frequent missions that NASA is planning for the early part of the next millennium.

7. ACKNOWLEDGMENT

The research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

8. REFERENCES

- Wyatt, E. J., et. al: 1997, "An Overview of the Beacon Monitor Operations Technology", *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, Tokyo, Japan.
- Wyatt, E. J., et. al: 1998, "Beacon Monitor Operations on the Deep Space One Mission", *Fifth International Symposium on Space Mission Operations and Ground Data Systems*, Tokyo, Japan.
- DeCoste, D.: 1997, "Automated Learning and Monitoring of Limit Functions", *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, Tokyo, Japan.
- Sherwood, R., et. al: 1997, "Flight Software Implementation of the Beacon Monitor Experiment On the NASA New Millennium Deep Space 1 (DS1) Mission", *Second International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations*, Oxfordshire, UK.
- Chien, S., et. al: 1997, "Resource Scheduling for a Network of Communications Antennas", *Proceedings of the IEEE Aerospace Conference*, Aspen, CO.
- Stahle, R.L., et. al: 1996, "Pluto Express: Advanced Technologies Enable Lower Cost Missions to the Outer Solar System and Beyond," *International Low Cost Planetary Missions Conference*, Laurel, MD.
- Wyatt, E.J. and J.B. Carraway: 1995, "Beacon Monitoring Approach to Spacecraft Mission Operations" *First International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations*, Oxfordshire, UK.