

Anomaly Trends for Long-Life Robotic Spacecraft

Nelson W. Green,* Alan R. Hoffman,† and Henry B. Garrett‡

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109

Three unmanned planetary spacecraft to the outer planets have been controlled and operated successfully in space for an accumulated total of 68 years. The Voyager 1 and 2 spacecraft each have been in space for more than 27 years. The Galileo spacecraft was in space for 14 years, including eight years in orbit about Jupiter. During the flight operations for these missions, a total of 3300 anomalies for the ground data and the flight systems have been tracked using the Jet Propulsion Laboratory's anomaly reporting tool. Methods and results are described for classifying and identifying trends relative to ground system vs flight system, software vs hardware, and corrective actions. Several lessons learned from these assessments can significantly benefit the design and planning for long-life missions of the future. These include the necessity for redundancy to ensure successful operation of the spacecraft, awareness that anomaly reporting is dependent on mission activity not the age of the spacecraft, and the need for having a program to maintain and transfer operation knowledge and tools to replacement flight team members.

I. Introduction

OPERATING long-life interplanetary spacecraft involves the monitoring of spacecraft functionality for both nominal and anomalous behavior. During flight operations for the Voyager and Galileo spacecraft, a total of 3300 anomalies has been tracked using the Jet Propulsion Laboratory's anomaly reporting tool for ground data and flight systems operations. These anomalous events require evaluation for the proper response required to ensure the continuation of the spacecraft's mission.

For this study, a long-life spacecraft is defined as one designed to function reliably for 10 or more years in the space environment. Although not specifically designed as such, the Voyager and Galileo spacecraft can be viewed as prototypes of long-life spacecraft because of their successful operation spanning several decades. The two Voyager spacecraft launched in 1977 have successfully completed flybys of several outer planets and now are leaving our solar system and flying trajectories taking them into interstellar space. Based on current consumables usage and continued mission operations attention, the spacecraft are expected to continue to return data until 2020. The Galileo orbiter was launched in 1989 and orbited Jupiter from 1995 to 2003 when it was intentionally impacted into the atmosphere of Jupiter just before its consumables were depleted.

Anomalies that occurred on these three long-life spacecraft have been analyzed with the intent of providing future designers and spacecraft operators an overview of when anomalies are most likely to occur, what kinds are most common, and what has historically been the response taken to correct for these anomalies. Studies of anomalies on Earth-orbiting spacecraft have been performed with the work of Sperber^{1,2} notable for its similarity to this study for its tracking of anomalies by time and subsystem. Others who have examined anomalies for interplanetary missions have concentrated on particular types of anomalies. Lutz and Mikulski³ studied

safety-critical software anomalies (i.e., anomalies that "posed serious threats to the embedded software systems"), whereas Brown⁴ and Kobele⁵ focused on the need for and use of spacecraft redundant systems. This study also examines the use of redundancies briefly as it illustrates a planned corrective action built into spacecraft to ensure the continuation of a mission in spite of an anomaly in a mission critical spacecraft subsystem.

II. Mission and Technology Descriptions

The technologies used on the Voyager and Galileo missions were state of the art when the spacecraft were designed, but if compared to modern standards, the electronics and technology would be considered obsolete. The Voyager missions, however, continue to transmit scientific and engineering data from deep space with technology more than 30 years old. These systems, and any made for future long-life missions, must be actively monitored and maintained by the flight team as the spacecraft continue on their extended missions.

Key features of the Voyager and Galileo spacecraft and missions are summarized in Table 1. Of particular interest are the limited onboard data storage and the slow communications data rate caused by hardware limitations. For Galileo, the loss of the high-gain antenna severely limited the rate of downlinking data. Primary power sources for these missions are radioisotope thermoelectric generators (RTG). Primary long-distance communications used both S and X bands using NASA's Deep Space Network receiver stations. Examples of the mission trajectories for Voyager and Galileo are given in Figs. 1a and 1b and are both examples of trajectories that utilized gravity-assist techniques during flybys of planetary bodies. A summary of the comprehensive environmental test program implemented for the Galileo spacecraft is given in the references.⁶

III. In-Flight Anomaly Assessment

For flight missions managed by the Jet Propulsion Laboratory (JPL), all in-flight anomalies are documented by incident, surprise, anomaly (ISA) reports. According to JPL documentation, "An Incident, Surprise, Anomaly (ISA) shall be written on all incidents, real or suspected, which indicate an anomaly in hardware, software, sequencing, test/operations procedures, etc." During flight operations to date for Voyager and Galileo (through mission end in September 2003), nearly 3300 ISAs have been generated. Most reports were processed in a hard-copy format and later transferred to an electronic database in the JPL Problem Reporting System. The assessments for this paper are based on the electronic versions of these reports and on earlier assessments of these missions.

Because of the large quantity of anomalies recorded, an in-depth assessment of all of the Voyager and Galileo ISAs was beyond the scope of this study. Galileo anomaly reports were examined in some

Received 8 November 2004; revision received 28 April 2005; accepted for publication 1 May 2005. Copyright © 2005 by the American Institute of Aeronautics and Astronautics, Inc. The U.S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for Governmental purposes. All other rights are reserved by the copyright owner. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0022-4650/06 \$10.00 in correspondence with the CCC.

*Associate Engineer, Mail Stop 185, Reliability Technology Group, 4800 Oak Grove Drive.


†Member, Engineering Staff, and Principal, Reliability Engineering Office, Mail Stop 122-107, 4800 Oak Grove Drive.

‡Member, Engineering Staff, and Principal, Reliability Engineering Office, Mail Stop 122-107, 4800 Oak Grove Drive. Associate Fellow AIAA.

Table 1 Spacecraft and mission information

Information	Voyager 1 and 2	Galileo orbiter
Spacecraft		
Number of science instruments	10	9 Orbiter/6 probe
Number of electronic parts	61,953	85,681
Power source	RTG ^a	RTG ^a
Number per spacecraft	Three	Two
Beginning of mission	480 W	570 W
Recent	320 W (February 2004)	435 W (end of mission)
Communication links	S-band, X-band (downlink only) 16–1400 bps	S-band, X-band (downlink only) 10–134.4 kbs (planned, 160 bps effective because of high-gain antenna failure)
Data storage	Tape recorder: 5.1×10^8 bits	Tape recorder: 9×10^8 bits
Mission		
Mission type	Flyby	Orbiter with probe
Primary mission design life	Through Saturn encounter	Five Jovian orbits
Solar distances design range	1–10 AU	0.6–5 AU
Launch year	1977 (both spacecraft)	1989
Key mission events	Jupiter: 1979 Saturn: 1980, 81 Uranus: 1986 Neptune: 1989	Venus: 1990 (gravity-assist flyby) Earth: 1990, 92 (gravity-assist flyby) Asteroids (Gaspia and Ida): 1991, 1993 Jupiter: 1995 through 2003
Current status (June 2004)	Voyager 1: 91 AU Voyager 2: 73 AU	Jupiter impact: 2003 (end of mission)

^aRTG-radioisotope thermoelectric generator.

**Fig. 1** Representative mission trajectories for outer-planet missions using gravity assist.

detail, but Voyager anomaly reports were only examined for numerical trends with time. In either case the intent was to develop information on anomaly types and occurrences to improve spacecraft design, reliability, and operations for long-life missions.


A. Anomalies vs Time

A look at anomaly reporting trends with respect to time for the Voyager and Galileo spacecraft provides a perspective for when anomalies are likely to be reported on other long-term missions. The trends reported here are not based exclusively on the types of technology used, though that is always an underlying factor, but rather on when in the mission were anomalies most often reported for these long-life missions. As expected, the largest number of reported anomalies came at the outset of each mission when spacecraft systems were first operated in flight and in the space environment. After one to two years of flight, anomalies of all types tended to decrease. This observation supports the hypothesis of the existence of a flight team "learning curve" during which mission controllers became used to flight operations and resolve initial anomalies. The number of reported anomalies also tended to increase just prior to and during major encounters. Both of these trends can be seen in


the ISA vs time plots for the Voyager and for Galileo spacecraft (Figs. 2a and 2b).

The Voyager spacecraft were launched in August and September 1977 with the primary mission of visiting both Jupiter and Saturn. In 1979 both of the spacecraft encountered Jupiter before continuing to Saturn in 1980 and 1981. After Saturn, Voyager 1 was directed out the ecliptic plane and began moving toward interstellar space. Voyager 2 continued on to visit Uranus in January 1986 and Neptune in August 1989 before heading out of the solar system.

For both spacecraft, the number of reported anomalies shows localized peaks at or near an encounter with a planet. (Note that the plotted ISA totals combine the reported anomalies for both Voyager 1 and 2 spacecraft.) Reasons for increases in reported events include uploading of maneuver software, refinements from the ground data system developed during preevent testing, and the reactivation of instruments that had been in a dormant mode during the cruise stage leading up to an encounter with a planet. In the latter case, any abnormalities or changes in the instruments caused by aging or space environmental effects would have been noticed upon reactivation and reported as anomalies in the JPL reporting system. After the Saturn encounters in 1980 and 1981, only one spacecraft, Voyager 2, continued on a trajectory to flyby Uranus and Neptune. The



a) Voyager combined ISA report quantities by year



b) Galileo ISA report quantities by year

Fig. 2 Total number of ISA reports per year for both the Voyager and Galileo missions. Note that ISAs for both Voyager 1 and 2 are included in the totals.


number of anomalies as a result of planetary encounter instrument reactivations decreased because they came from a single spacecraft.

The Galileo plot of anomalies as a function of time (Fig. 2b) shows a large number of reported anomalies in the early stages of the mission. These reported anomalies stem from both the initial operation of the spacecraft and early encounters with planets for gravity assists. To arrive at Jupiter and its moons in December of 1995, the spacecraft trajectory required gravity-assist encounters with Venus and Earth in 1990 and a second gravity-assist encounter with Earth in 1992. There were also asteroid flybys in 1991 and 1993 involving the activation of at least part of the spacecraft instrumentation.


Reported anomalies peaked in 1990 corresponding to both the first year of operation and two planetary encounters. Another peak in anomaly reports occurred once Galileo reached Jupiter and entered orbit around the planet. Galileo had a total of 35 encounters with the planet and its moons rather than the two planetary encounters of Voyager 1 and the four by Voyager 2. The plot of anomalies for Galileo shows a slight increase in activity in 1995 as the craft was readied for insertion into orbit around Jupiter and a large increase in anomalies in 1996 corresponding to the first set of Jupiter moon encounters. A more detailed look at these two peak anomaly periods for Galileo is given in Fig. 3.

Anomaly report quantities have been plotted for two significant portions of the Galileo mission. The first 18 months after launch included the initial in-flight learning period for the mission team and the first two gravity-assist planetary flybys. On a scale showing ISA report quantities by month (Fig. 3a), the trend of relatively large numbers of reports at the beginning of the mission and near major encounters is clearly seen. In the time between flying by Venus and the first Earth flyby, the number of anomaly reports dropped significantly. Additionally, the number of ISA reports dropped off again after the Earth flyby in December 1990. Although it is not shown in this plot, there is a similar decrease and increase before the second Earth flyby in December 1992.

The second detailed plot (Fig. 3b) shows the primary mission for the Galileo orbiter once it had reached Jupiter orbit in late 1995. Surprisingly there was not a rise in anomalies before insertion of the spacecraft into Jovian orbit in December 1995, but there was a sharp increase prior the first encounter with a Jovian moon with the



a) ISA quantities for the first 18 months of flight



b) Monthly ISA quantities for the primary mission

Fig. 3 Number of ISA reports for the Galileo mission in month long increments for a) the first 18 months after launch and b) the three years of the primary mission. Peaks in the number of reports are seen immediately after launch, during planetary flybys of Venus and Earth, and the first few encounters with the moons of Jupiter: JOI, Jupiter orbit insertion; Jovian moon encounters are listed by the first letter of the target moon and the orbit number; E4 targeted Europa and was the fourth orbit of the planet.


Ganymede flyby at the end of June 1996. The number of anomaly reports decreased from this first moon encounter, but increases are generally found in proximity to moon flybys.

What is surprising about all of the plots showing the number of ISAs as a function of time is that the number of anomalies reported for all spacecraft sharply declines with increasing time. This trend is in direct opposition of the expectation that the spacecraft will become less and less operational with time as the devices on board age and absorb more planetary and cosmic radiation. The decrease in anomaly reports can be related to fewer active devices onboard the spacecraft (e.g., encounter science instruments and maneuver controls). Instruments are turned off either as a result of inactivity, as in the case of the Voyagers, or inoperability caused by age or radiation exposure.


B. Anomalies vs Workforce

The human side of the anomaly reporting system must also be examined when plotting the numbers of anomalies through the life cycle of a mission. It might seem obvious, but without a mission team noticing and reporting them anomalies will occur and not be reported or recorded. Plots showing the number of ISA reports and total workforce employed on the mission team (Fig. 4) show that there is a fair correlation between the size of the workforce for a mission and the number of anomaly reports generated. The total workforce includes both employees and contractors and is given for each fiscal year as opposed to the calendar year used for the ISA totals.

The Voyager missions (Fig. 4a) show a fairly strong relationship between the workforce total and the number of anomalies reported with peaks at the mission start and with each encounter thereafter. The workforce totals show a diminishing trend overall, but drop significantly after the last Voyager encounter with Neptune in 1989



a) Voyager combined ISA and workforce totals



b) Galileo ISA and workforce totals

Fig. 4 ISA report totals plotted along with the total workforce for both the Voyager and Galileo missions. Note the fairly close tracking of workforce totals with the ISA totals for the Voyager mission workforce. Values for both Voyager 1 and 2 are combined in this plot. The Galileo plots show much less of a correlation between the workforce and ISA report quantities. Much of the difference at the beginning of the mission was caused by the failure of the high-gain antenna and the work needed to design a work around solution and the subsequent reprogramming of the flight software. After the arrival at Jupiter, the workforce correlates more closely with the number of ISA reports.

with a corresponding drop in the number of anomalies reported for both spacecraft. When examining these workforce plots, it is important to keep in mind that the number of personnel on a project, and potentially the number of reported anomalies, can be related to the decreasing amount of spacecraft maneuvers required when the spacecraft is in a steady-state cruise configuration as well as the shutdown of some of the science instruments onboard both of the Voyager spacecraft. This has been particularly true of the Voyager spacecraft as they fly into interstellar space.

The time plot of anomalies reported along with the workforce totals for the Galileo mission (Fig. 4b) shows far less correlation between the two plots than could be seen in the plot for the Voyager missions. The divergence of the number of anomalies and the number of personnel working on the project is caused by the early difficulties encountered by the Galileo mission (i.e., the high-gain antenna deployment difficulties in April 1991 and postlaunch ac/dc bus imbalances). To understand and work around the problems encountered in flight, the workforce increased during the cruise stage of the mission and only decreased after Jupiter orbit insertion (JOI) in late 1995. Specifically, 80% of the ~70,000 lines of code for the onboard software in the command data subsystem had to be modified prior to JOI because of the new mission plan to perform the mission without a high-gain antenna. Subsequently for code maintenance and corrections only 2% had to be revised. After JOI, the workforce totals and number of reported anomalies follow more closely together through the end of the mission in 2003.

As with the plot for the Voyager missions, the number of anomalies reported for the Galileo mission and the workforce totals decreased with time especially after a major encounter. With the Galileo mission, however, the spacecraft did not leave the vicinity of Jupiter but continued to make encounters with the planet and its several moons on a regular basis with periods of nonactivity between each encounter. It is therefore interesting that the number

of anomalies followed the downward trend with time seen on the Voyager missions in concert with a decrease in total workforce.

C. Sources and Corrective Actions

For the Galileo mission, a closer look was taken at the types of anomalies reported in the ISAs and the corrective actions used to resolve them. Each ISA was reviewed and categorized by both the source of the anomaly and the corrective action taken to address the anomaly. The anomaly sources were intentionally simplified to procedure, ground hardware, flight hardware, ground software, and flight software. Likewise, the corrective actions taken were categorized as undetermined, use as is, ground hardware, ground procedure, flight procedure, ground software, and flight software. Pie charts showing the percentages of anomalies from each source and the corrective actions taken in response to the ISA reports are given in Figs. 5a and 5b.

Over the span of the entire mission, flight hardware systems were the most common origination point for reported anomalies. Flight hardware anomalies, defined for this study as being any unexpected changes or incidents related to the physical devices on the spacecraft, represented 41% of all anomalies. Software-related anomalies represented a total of 42% of the anomalies reported, but these were divided into two sections: flight and ground software issues. Flight software anomalies, 31%, were defined as incidents involving the code running on the spacecraft. All other software-related anomalies were classified as originating with ground software and represented the other 11% of the software anomalies. There was some overlap in this category because some ground software was designed to create flight software but would still be considered ground software for this study because it operated on ground-based computers. Hardware used at ground stations, including both mission control and receiving stations, originated 6% of the ISA reports, whereas procedure issues, including mission rules and operating guidelines, generated the remaining 11% of all anomaly reports.

Corrective actions were also divided into primary categories for the purposes of this study. Corrective actions, in general, are the






Fig. 5 Pie charts showing a) the sources of anomaly reports and b) the corrective actions taken as a percentage of all Galileo ISA reports. The largest sources for anomalies in a) are flight hardware and software (flight and ground combined). The most common corrective action taken in b) is use as is followed by software (flight and ground combined) and a procedural change.

Table 2 Anomaly sources with their most common corrective actions

Anomaly source	Corrective action taken
Flight hardware	Use as is: 51%
	Flight software: 25%
	Flight procedure: 15%
Flight software	Use as is: 32%
	Flight software: 26%
	Ground software: 19%
	Flight procedure: 16%
Ground hardware	Ground hardware: 57%
	Use as is: 22%
	Ground procedure: 16%
Ground software	Ground software: 51%
	Flight procedures: 23%
	Use as is: 11%
Procedure	Flight procedures: 42%
	Ground procedure: 37%
	Use as is: 10%

**Fig. 6** Plot of actions taken to correct anomalies reported in ISA reports as a function of time. Note that the most common response throughout all phases of the mission is use as is.

responses taken by the mission team to the anomalies recorded in the ISA reports. Over the scope of the entire mission, the decision to take no corrective action, or as stated in the corrective action field of the ISA report, to use as is, was consistently the most common response. This is true both when examining the anomalies from each year (Fig. 6) and when looking at the corrective action totals over the life of the mission (Fig. 5b). In the latter case, it can be seen that it was decided to take no action for 35% of all anomalies. As expected, use as is was most commonly used when the system originating the anomaly was flight hardware related as seen in Table 2, where the top corrective action for each anomaly source is shown. Some flight hardware anomalies were addressable with software corrective actions, but in many cases no action could be taken either because the system was unfixable via software or the anomaly was not mission threatening. The use-as-is response was also commonly used for anomalies that happened only once or for incidents that were considered to really reflect normal behavior (i.e., an incident that was surprising at first but was determined to be within normal mission parameters upon closer examination). Of particular interest is that in 1996, the first year after reaching Jupiter when instruments were reactivated and first put to use in the Jupiter environment, the number of anomalies increased significantly, and the most common response to the incidents was to use as is.

When an anomaly could be addressed by corrective action, the most common action was a software modification. When looking at the timeline of the Galileo mission, software corrective actions peaked at the beginning of the mission and decreased with time with a slight increase following the arrival at Jupiter. For this analysis, software corrective actions were divided into two types of corrective action, those involving software transmitted to the computers flying on the Galileo spacecraft and actions addressing software used on ground-based computers. Of the two, changes to flight software were more numerous. Together software corrective actions almost equal to the number of use-as-is responses to anomalies and represent another third of the total corrective actions.

Flight software corrective actions were the most common responses to anomalies found in both flight software and the hardware on the spacecraft. Typical corrective actions of this type included powering on or off particular instruments or subsystems or enacting a workaround procedure so that the mission could continue. Other anomalies that required a flight software corrective action commonly originated in previous versions of transmitted flight software. Some of the corrective actions involving flight software were incorporated into scheduled software updates, whereas others were done in real time in response to an anomaly that needed to be addressed immediately.

Corrective actions involving changing software for ground-based computers were fewer in number than those transmitted to the Galileo spacecraft, but sometimes they intersected with flight software issues. Some of the software corrective actions were to ground software that created the sequences eventually transmitted to the spacecraft. This kind of fix was counted as a ground software fix, but it directly related to software transmitted to the spacecraft. Other ground software corrective actions addressed anomalies found in software in use in ground-based computers. A small amount of ground software corrective actions was made to work with flight hardware-based anomalies that could not otherwise be addressed.

The third most common type of corrective actions was a procedural change, either for procedures related to flight operations or for those dealing with ground operations. Flight procedural corrective actions, that is, revisions to flight mission rules used to determine spacecraft operations, were most common in the early part of the Galileo mission and in 1996 following Galileo's arrival at Jupiter. It is supposed that procedural changes were at their peak during these times because the mission team was learning how the spacecraft operated in flight. This supposition is supported by the decrease in flight procedural corrective actions with time after the first two years of flight and after the arrival at Jupiter.

All other procedural corrective actions were considered to be ground based. These corrective actions follow a similar trend to flight procedural corrections in that they were more numerous in the beginning stages of the mission. They differ, however, in that they do not seem to be as closely correlated to Galileo's arrival at Jupiter.

The two additional types of corrective actions determined in this study were those related to ground hardware and those here classified as undetermined. Ground hardware corrective actions involved changing or modifying ground support equipment either in Mission Control or at one of the Deep Space Network locations and were generally made in response to some ground or flight hardware need. Undetermined corrective actions were those where the electronic version of the ISA was unclear regarding the type of corrective action taken or whose corrective actions were detailed in attachment files missing from the database.

IV. Redundancy

Part of the planning for a long-life mission is how to address anomalies occurring during the mission that cannot be corrected remotely. In many cases, particularly for critical systems, this means adding redundant systems to the spacecraft as discussed in depth by Kobele.⁵ In the early design phase for spacecraft, the design team must work with hardware designers and mission planners to perform a series of cost vs risk studies to determine the benefits of including redundant hardware elements in the spacecraft design. The designers must evaluate the criticality of the hardware for mission success, the availability of redundant elements, and the estimated reliability of the system. Opposing the addition of redundancies are increased mass, additional power usage, the need for additional hardware and software to support the inclusion of the hardware, and the increase in funding needed to supply the redundancy.

Redundancy for a flight spacecraft can be achieved by the following methods: block (simple) and functional. Block redundancy consists of a duplicated hardware set that replaces a failed unit when a failure is detected. Functional redundancy consists of replacing performance functions by utilizing performance aspects of other subsystems. As an example, functional redundancy was used for

Table 3 Voyager and Galileo in-flight failures salvaged by redundancy

Spacecraft	Failure description	Subsystem	Cause	Time of occurrence	Redundancy applied
Voyager 1	X band traveling-wave tube degradation	Radio frequency subsystem	Unknown, possibly random aging	10.2 years	Block X-band traveling-wave tube swap by ground command
	Lost S band downlink	Radio frequency subsystem	Component failure: ultrastable oscillator; possibly random aging	15 years	Block Automatic exciter swap and auxiliary oscillator in new exciter used
	Rapid degradation in yaw limit cycle	Propulsion subsystem	Thruster plugged	22 years	Block Auto swap to redundant branch
Voyager 2	Pyro amps "A" missing at RTG boom release	Power/pyro subsystem	Unknown	0 year (at launch)	Block Pyro circuit is inherently redundant
	Receiver 1 failed 20 min after turn on	Radio frequency subsystem	Hardware design	7.5 months	Block
	Lost S/C data when memory "B" block 256 memory failed	Flight data subsystem	Part failure	4.1 years	Block Switch to memory A for rest of mission and S/W upload changes
Galileo	High-gain antenna failed to deploy	Antenna subsystem	Unknown	17.8 months	Functional Low-gain antenna and operational workaround

the Galileo high-gain antenna because including a redundant large-diameter antenna was not practical. The fault-tolerant design for this subsystem used a low-gain omnidirectional antenna at a reduced data rate as a degraded, but acceptable, redundant system. In this case, the redundant system is not a direct duplicate of the original system but provided the same function even if in a reduced capacity.

As previously determined by Brown,⁴ the use of redundancy on both Voyager and Galileo was essential for the completion of their missions. The use of block and functional redundancy for these missions is summarized in Table 3. These redundant systems have been used during all phases of the missions: launch phase, cruise phase, encounter phase, and extended mission phase.

As part of the Voyager and Galileo design teams tradeoff studies relating to redundancy, the flight histories of the Mariner and Viking robotic spacecraft of the 1960s and 1970s were examined. Decisions were made to include redundancy in critical engineering subsystems, power conditioning and distribution, communications, and spacecraft attitude control and stabilization. Brown evaluated the usage of redundancy for unmanned planetary missions and concluded that block redundancy was necessary for a mission to meet the science objectives.

Of particular importance to both the Voyager and Galileo missions have been redundancies in the communication systems. All three spacecraft have experienced anomalies and failures in the communications equipment that required the use of a redundant system. Without the redundant system the spacecraft would have been unable to communicate to the mission team with the effective loss of the mission.

Specific examples of the use of redundant communication systems include the failure of one of Voyager 2's receivers 20 min after turn on and two failures on Voyager 1: degradation of an X-band traveling wave tube and failure of an S-band downlink ultrastable oscillator channel. In all cases, hardware (block) redundancy was available and used through either automatic switching or through a command sent to the spacecraft from the mission team. These failures occurred at a range of times after launch, 7.5 months for the Voyager 2 failure, and 10 and 15 years, respectively, for the two Voyager 1 communication system failures.

The failure of deployment of the Galileo high-gain antenna a year and half after launch resulted in a major effort by the flight team and industry experts to rectify the problem without the use of redundant systems. Several techniques to dislodge the stuck antenna were attempted including turning the spacecraft to heat the ribs and mechanical "banging." None of these worked. Because of the size of the antenna (3.7 m in diameter), it was impractical to fly a redundant antenna. The only backup system for the high-gain antenna was a low-gain omnidirectional antenna whose primary function was to

permit communication with the spacecraft during the early stages of the flight and during spacecraft maneuvers. After significant study, the decision was made to use functional redundancy that utilized the onboard capabilities for the recording of scientific data on the tape recorder for subsequent transmission to the Earth ground stations at low data rates. By reconfiguring the ground stations, the effective received data rates were increased. As a result of the use of functional redundancy and the operational workarounds onboard the spacecraft and on the ground, more than 70% the original science objectives were achieved with a significant increase to the scientific knowledge of Jupiter and its moons.

The monitoring of engineering data, mission team reaction, and timely assessment are key aspects that need to be designed into long-lived spacecraft missions. In this regard retaining robustness in communication links is of primary importance. Onboard autonomous swapping of critical subsystems must also be built into the architecture of the spacecraft system when communication links are many hours long because of the distance between the spacecraft and Earth. Redundancy and its usage must be evaluated in the design and resources trades that occur for long-life missions. Catastrophic failures would have been the outcome for the Voyager and Galileo missions if redundancy had not been available.

V. Knowledge Transfer

Designing, building, testing, and flying unmanned robotic scientific spacecraft require the combined efforts of many technical specialists working over a period of several years. Few of these individuals will stay with the project throughout the entire scope of the project, especially if it is to be a long-life mission. Each specialist will contribute to their portion of the project, design a system, develop software, test a subsystem, or operate an instrument before being reassigned to another project where their expertise is required. This division of labor is appropriate because of the nature of the changing needs of a project throughout its life cycle, but it requires an effective means of transferring knowledge about the different systems of the spacecraft so that information is not lost when the original specialist leaves the project. Knowledge transfer can be either through training of personnel or, more appropriately, through good documentation so that work done by the different groups of specialists can be reviewed and understood at a later time. This last means of knowledge transfer is particularly necessary for long-life missions as the availability of the original specialists will decrease with time because of a career change, retirement, or death.

In 1999, 22 years after the launch of Voyager 1, the flight team observed that one of the thrusters was plugged on the primary loop for the attitude control subsystem. Because the switchover would involve writing some new flight code and changing some flight

parameters, a retired software specialist familiar with the attitude control software onboard the spacecraft was tasked with helping develop and review the software upload. The switch was successfully performed, and the onboard software maintains the required three-axis stabilization of the spacecraft. In this case, an original specialist was located and asked to design an appropriate corrective action for an anomaly on the spacecraft, but if the specialist had not been available the mission team would have needed to develop the software fix from existing documentation. In some ways the decision to bring in an original specialist for anomaly resolution shows a concern about the transfer of knowledge.

Effective knowledge transfer is difficult to achieve. Documentation traditionally covers only the basics of a system or software package. Transferring root knowledge and underlying assumptions held by the system or software designers is a much more difficult task. But for long-life spacecraft where system failure might need to be addressed long after the development of the spacecraft, these kinds of knowledge are still needed. Mission teams need to be supported by specialists who know how to work with the original systems and software, no matter how old and obsolete they become.

VI. Lessons Learned

The lessons learned from the assessments of the flight anomalies that have occurred during the accumulated flight time of 66 years for unmanned outer planets mission are as follows:

- 1) The number of anomalies is dependent on mission activity with peaks occurring during launch and early cruise, preencounter testing, and during an encounter.
- 2) The corrective action most frequently noted was "use as is," the second was a software update, and the third involved changes to procedure.
- 3) Block and/or functional redundancy have been necessary for the successful operation of the spacecraft.
- 4) Robustness in the underlying architecture of the system design has to be built in especially for communication systems.
- 5) Resource planners for future long-life missions must provide knowledge management for the operations staff in either training or in detailed documentation so that anomalies can be effectively resolved.

VII. Summary

Anomaly reports, specifically Jet Propulsion Laboratory's postlaunch incident surprise and anomalies (ISA) reports, have been analyzed for the Voyager and Galileo deep-space spacecraft with an accumulated flying time of more than 66 years. The number of anomalies noted for these spacecraft display peaks corresponding with increased mission activity. Particular evidence of these increases in reported anomalies occurred immediately after launch

and in proximity to major encounters with planets and moons. In the case of Galileo, anomalies were equally spread between flight hardware and software. Corrective actions taken in response to anomalies were most commonly either use as is or a change in software transmitted to the spacecraft. Functional or block redundancy was used on all of these missions to overcome failures and continue to provide useful data to the science teams with redundancy in the communication system being most critical. Such redundancy should be required for future long-life missions and must be carefully selected and analyzed during prelaunch development. Although rigorously analyzed before launch, flight teams must continue to monitor spacecraft health and evaluate trend data throughout the mission life because mission critical anomalies can occur at any time in a mission from immediately postlaunch to after decades of flight. To effectively deal with anomalies, knowledge about the systems and software must be transferred as personnel change over the lifecycle of the mission. Key areas of knowledge transfer include skill retention in progressively obsolete systems and proper documentation for systems and instrument control.

Acknowledgments

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA. The authors gratefully acknowledge the following colleagues for their contributions of data and expertise: T. Brady, A. Brown, R. Draper, T. Hogle, S. James, R. Jones, S. Krasner, and N. Palmer.

References

- ¹Sperber, B. R., "The Public Record of Commercial Satellite Anomalies vs. Time," *CERT'91: 2nd Components Engineering, Reliability and Test Conference*, Electronic Components Inst. Internationale, Ltd., Crowborough, East Sussex, England, U.K., 1991, pp. 101–106.
- ²Sperber, B. R., "Better with Age and Experience—Observed Satellite in-Orbit Anomaly Rates," AIAA Paper 94-1083, Feb.–March 1994.
- ³Lutz, R. R., and Mikulski, I. C., "Empirical Analysis of Safety-Critical Anomalies During Operations," *IEEE Transactions on Software Engineering*, Vol. 30, No. 3, 2004, pp. 172–180.
- ⁴Brown, A., "Flight Anomaly Characterization as a Means of Evaluating Product Assurance Requirements," *Proceedings of the 15th Aerospace Testing Seminar*, Inst. of Environmental Sciences, Manhattan Beach, CA, 1994, pp. 289–296.
- ⁵Kobe, P., "Maintainability of Unmanned Planetary Spacecraft: A JPL Perspective," AIAA Paper 89-5070, July 1989.
- ⁶Hoffman, A. R., "Galileo Environmental Test and Analysis Program Summary," *Proceedings of the 13th Aerospace Testing Seminar*, Inst. of Environmental Sciences, Manhattan Beach, CA, 1991, pp. 185–202.

J. Korte
Associate Editor