

Fault Protection Techniques in JPL Spacecraft

Paula S. Morgan

Jet Propulsion Laboratory/California Institute of Technology, Pasadena, California
(818-393-1092) Paula.S.Morgan@jpl.nasa.gov

Abstract: For all JPL spacecraft, maintaining the health and functionality of spacecraft subsystems and science instruments is an ongoing task; a challenge which must be met throughout the lifetime of every mission. Material stresses in flight caused by solar heating, the cold of deep space, solar radiation bombardment, etc., can degrade the mission or contribute to malfunctions in subsystem components. In addition to these health risks, flight software sequences and coding updates periodically sent to the spacecraft can potentially introduce human error. As spacecraft design sophistication and complexity increases, fault diagnosis and resolution becomes a more difficult and time-consuming task for the Spacecraft Operations Ground-based Team who must collect large volumes of telemetry data to diagnose faults. These telemetry streams contain hundreds of system data products that must be compared to archived historical data and spacecraft design information to determine fault causes and resolution actions. Additionally, those spacecraft missions that experience great Earth-spacecraft distances (such as outer planet exploration) present an additional challenge, as the ever-increasing delay period between commands sent and received by the spacecraft limits the ability to respond to fault occurrences in a timely manner. Time delays also present problems when spacecraft mission objectives contain “crucial events” that must take place at specific times, or when serious, potentially mission-catastrophic faults must be fixed immediately.

Fault management may be approached by implementing functional redundancy, redundant hardware, and Fault Protection (FP) techniques. This strategy provides autonomous monitoring of component operation, device health, internal & external conditions, and power allocation, by responding to any anomalous conditions through automated responses containing preprogrammed instructions. Thus, mission integrity may be optimized by implementing Fault Protection strategies which will provide a more robust spacecraft system with greater diagnostic capabilities.

While every JPL spacecraft requires some unique mission specific fault protection, there are many requirements that are common to all spacecraft configurations. These consist of protecting command and data processing & attitude control computers, protection against communication loss with the spacecraft, ensuring that safe external and internal temperature levels are maintained, and recovery from power overloads. Additionally, most JPL spacecraft are equipped with a general-purpose “Safe Mode” response that configures the spacecraft to a lower power state that is safe and predictable so that diagnosis of more complex faults can be addressed by the Operations Team. This paper details the generic application of fault protection techniques which are implemented into most JPL spacecraft designs.

INTRODUCTION

Once JPL spacecraft are ferried out of Earth's gravity well, usually by multi-stage rocket, it will either enter Earth's orbit or proceed out into deep space. Through the use of NASA's Deep Space Network radio telescope system, the spacecraft's Operations Ground-based Team will stay in contact with the spacecraft, providing instructions through “uplinked” commands while the spacecraft's “downlink” telemetry stream provides detailed information of what it encounters throughout its mission. Once all systems are deployed, configured, and verified working after launch, the propulsion system will be utilized to target the spacecraft to the intended destination through its trajectory. For JPL's interplanetary spacecraft missions, objectives consist of orbiting or flying by an object, moon, or planet, or landing the spacecraft or its probe on the object it is encountering. The suite of scientific instruments carried on board the spacecraft will perform scientific tasks, some perform evaluations throughout the lifetime of the mission.

As spacecraft make their journey through the vastness of space, there are many influences that will provide a challenge in maintaining spacecraft health and functionality. All of these risk factors must be taken into account when designing JPL spacecraft, whose resolution may be facilitated substantially by implementing automated fault protection techniques.

1.0 HEALTH & SAFETY CONCERNS FOR DEEP SPACE MISSIONS

For spacecraft to function properly, external and internal influences must be monitored, regulated, and controlled during the entire lifetime of the mission. One of the most detrimental external influences on spacecraft operation in the vacuum of space is exposure to the Sun when the vehicle is in close proximity to this celestial body. Spacecraft surfaces superheat when exposed to the Sun, while shadowed surfaces can fall to extremely low temperatures. Material stress can result from this thermal expansion-contraction. This uneven heating can lead to warpage, camera distortion, or breakage of components. Some spacecraft will be equipped with fault-preventative devices to help alleviate some of these problems such as optical solar reflectors, mirror tiles, or multi-layer insulation thermal blankets which will reflect the Sun's heat and radiation so that the spacecraft is somewhat protected against overheating, while retaining its internal heat to prevent too much cooling. But fault protection techniques are also required to prevent an adverse thermal environment, as computers and spacecraft components will cease working if spacecraft temperatures become too extreme (Qualitative Reasoning Group 2005, DS-1 Project).

Precautions must also be taken to ensure that instruments do not fall out of operating limits, since many devices will only operate within a narrow range of temperatures. The spacecraft's interior environment must also be properly managed as well, as heat build-up can occur from the spacecraft's own systems. One method used to regulate internal temperatures is the circulation of spacecraft's gas or liquids (fuel) to cool its interior, so that the thermal state of these substances must be maintained so they do not freeze from deep space exposure. This condition would also render the propellant unusable so that the spacecraft would not be able to maneuver, eventually becoming misaligned with Earth so that no signals could be sent or received by the spacecraft.

Another source of error is human interaction with the spacecraft. Although precautions are taken to prevent the possibility of human-induced electro-static discharge events within spacecraft components during the manufacturing process, "latent failures" can sometimes present themselves after launch, rendering the device useless or partially useless. Additionally, human error can occur within command sequences that are periodically sent to the spacecraft. These commands contain instructions to control the spacecraft's activities either immediately, or over an extended period of time. These tasks consist of activities that must be performed during flight such as tracking Earth, monitoring celestial references for attitude targeting, performing maneuvers to fine-tune the trajectory when required, and carrying out science calibration and operations. Unfortunately, these command sources are subject to errors that can potentially cause serious faults. An example would be accidentally turning off a radio transmission or receiver device on board the spacecraft, which can lead to an inability to communicate with the spacecraft. Another fault might be turning on too many components at the same time so that the spacecraft's power source (solar panels, Radioisotope Thermoelectric Generators (RTGs), fuel cells, etc.) is unable to provide the power required. Such a condition will create a spacecraft-wide "under-voltage power-outage" to occur, in which loss of power to critical devices such as the computers, which must maintain their power levels to retain computer memory, can result (JPL 2005, Basics of Space Flight).

In dealing with these health and safety concerns, there are several limiting factors that require consideration. One of these is the ever-increasing lag time experienced on large Earth-spacecraft distance missions. Even though radio waves travel at the speed of light making spacecraft/Earth transactions almost instantaneous near Earth, as the distance between Earth and the spacecraft increases, even a signal traveling at the speed of light can take hours. This lag time becomes a high-risk deterrent to fault recovery when spacecraft are sent out great distances. Under some anomalous conditions, it is impossible for spacecraft to respond to ground commands in time to preclude a catastrophic failure. An example would be failure of a latch valve to close properly after re-pressurization of the spacecraft's fuel tanks. This fault could cause the pressure in the tanks to rise substantially in a very short period of time. If this condition occurred on the Cassini spacecraft for example, the pressure level could potentially reach a catastrophic level before the

pressure measurement data can even reach Earth to indicate the fault condition, since Cassini's signals take well over an hour to reach Earth from its Saturn-Titan orbit. Also, lag time becomes a significant factor for many spacecraft missions that contain one-time opportunities such as planet/moon encounters. For these events, the timing is crucial since only one opportunity exists to meet the objective and there may be no second chance. These unique events must proceed without fault interference in order for the spacecraft's mission to be successful.

As spacecraft design becomes more complex, fault diagnosis and resolution becomes a more difficult and time-consuming task to undertake. Fault causes can lead to a plethora of possibilities, which poses a substantial challenge for the Operations Team who must collect large volumes of telemetry data to diagnose faults and propose resolution actions. This manual process requires that hundreds of data products from the spacecraft's telemetry stream be compared to archived historical data and design information to evaluate the problem and propose a solution.

Hence, the common problems experienced by most spacecraft may be alleviated by implementing autonomous solutions within the spacecraft itself through the application of fault protection techniques.

2.0 FAULT PROTECTION IMPLEMENTATION APPROACH

Although each JPL spacecraft is unique in its configuration and mission objectives, the task of implementing autonomous fault protection may be approached in a generic manner. Some spacecraft designs may be quite simple (e.g. lack propulsion and attitude control subsystems entirely, such as an atmospheric probe), and some are quite complex but all spacecraft share common systems which require a similar approach in fault protection design.

Fault protection is applied by implementing functional redundancy, redundant hardware, and autonomous fault protection monitors and responses. Fault protection is used to facilitate redundant unit swaps, in addition to maintaining spacecraft health and safeguarding its operation through continuous monitoring of spacecraft systems. Anomalous conditions invoke fault responses that contain "preprogrammed instructions" to place the spacecraft in a safe state.

In general, autonomous fault protection should only be implemented on-board the spacecraft for those fault conditions where a ground response is not feasible or practical, or if fault resolution action is required within a pre-defined period of time of detecting the failure. Otherwise, the ground system should have adequate time to respond to the fault and should be responsible for the fault recovery. In both cases, the ground is responsible for failure diagnosis and re-configuration of the spacecraft to nominal operations after the fault.

2.1 Standard Fault Protection Implementation

Some spacecraft have design configurations simple enough to warrant only minimal fault protection which is meant to address any type of fault condition. Other spacecraft designs are very complex and sophisticated, have long mission durations, and must maintain a system with numerous error possibilities. Most spacecraft typically rely on a "general-purpose, Safe Mode" fault response which typically configures the spacecraft to a lower power state by powering off all nonessential spacecraft loads, commanding a thermally safe attitude, providing a safe state for the hardware, establishing an uplink and a downlink, reconfiguring to a low-gain antenna, and terminating the sequence currently executing on the spacecraft. This response is used to configure the spacecraft in safe and predictable state so that the Operations Team may have enough time to evaluate the fault causes and determine a solution.

Standard fault protection also includes an automated response to address "loss of spacecraft signal" faults that affect the Ground Operations Team's ability to communicate with the spacecraft. Failure to receive the spacecraft's uplink signal can be caused by a number of problems; ground antenna failures, environmental interferences, and spacecraft hardware failures, as well as an erroneous spacecraft attitude (pointing error), radio frequency interferences, or an error in an uplinked sequence (i.e. radio device accidentally turned off). If the spacecraft has experienced this type of failure and is no longer able to receive commands from the

ground, the fault protection response will attempt to re-establish the uplink. This type of fault protection is referred to as a “Command Loss Response” (from the perspective of the spacecraft, that it is no longer receiving ground commands) and is typically an “endless-loop” response.

Another type of standard fault protection is recovery from a system-wide loss of power. This is referred to as an “Under-Voltage” response, and could be caused by a number of faults depending on the spacecraft design (i.e. oversubscribing the power available, a short in the power system, or a communications bus overload). Should a system-wide power loss occur, not even the Safe Mode response will execute since the main computer will also lose power, causing loss of the mission. Fault protection must be implemented to sense the power drop so that the system may shed its non-essential loads from the communications bus, isolate the defective device, and re-establish essential hardware. The quick actions of this response allow critical spacecraft memories to be maintained throughout the Under-Voltage event.

Figure (1a) through Figure (1c) show three JPL spacecraft designs with quite different mission objectives, which employ most standard fault protection. Their mission design unique fault protection is also listed.

Figure (1a). *CloudSat Spacecraft FP Allocation*

CloudSat: *Earth Orbiting Satellite*

Standard FP: 3 Safe Mode Responses
5 Under-voltage Responses
Memory Scrubber & Bus FP

Unique FP: Significant computer & thermal FP

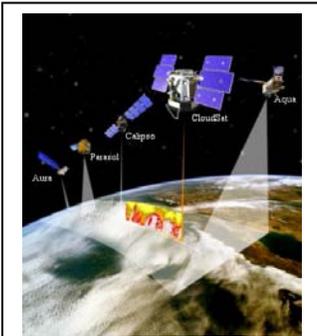


Figure (1b). *Stardust Spacecraft FP Allocation*

Stardust Spacecraft: *Inner Solar System; Comet Explorer*

Standard FP: 1 Safe Mode Response
1 Under-voltage Response
1 Command Loss Response
Memory Scrubber & Bus FP

Unique FP: Some computer & thermal FP

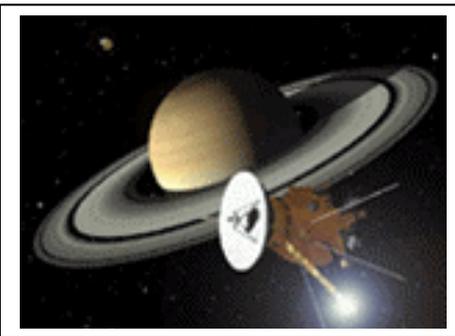


Figure (1c). *Cassini Spacecraft FP Allocation*

Cassini Spacecraft: *Outer Solar System; Saturn-Titan Explorer*

Standard FP: 1 Safe Mode Response
1 Post-Safe Mode Response
1 Under-voltage Response
1 Command Loss Response
Memory Scrubber & Bus FP

Unique FP: Significant command & data processing computer FP, radio unit FP, thermal FP, fuel tank pressure FP, attitude articulation and control computer FP



3.0 FAULT PROTECTION APPLICATION

Acknowledgement: The author would like to acknowledge following section from Reference: CAS-3-330 Fault Protection Requirements, Cassini Project, which was written by Sarah Gavit, Jet Propulsion Laboratory California Institute of Technology.

Fault protection responsibility is allocated to both Ground Operations Team and the Spacecraft. The spacecraft must deliver sufficient information on system health and fault recovery actions to facilitate spacecraft recovery by the Ground Team or the automated fault protection. On the spacecraft, autonomous fault protection is divided into two applications: either "Subsystem Internal Fault Protection" (SIFP), which is localized to subsystem components, or "System Fault Protection" (SFP), which will monitor and address faults that affect the entire spacecraft. Fault protection is allocated to SIFP design if the subsystem can recover itself without affecting the functionality or standard operation of another subsystem. The diagram below details this generic design approach:

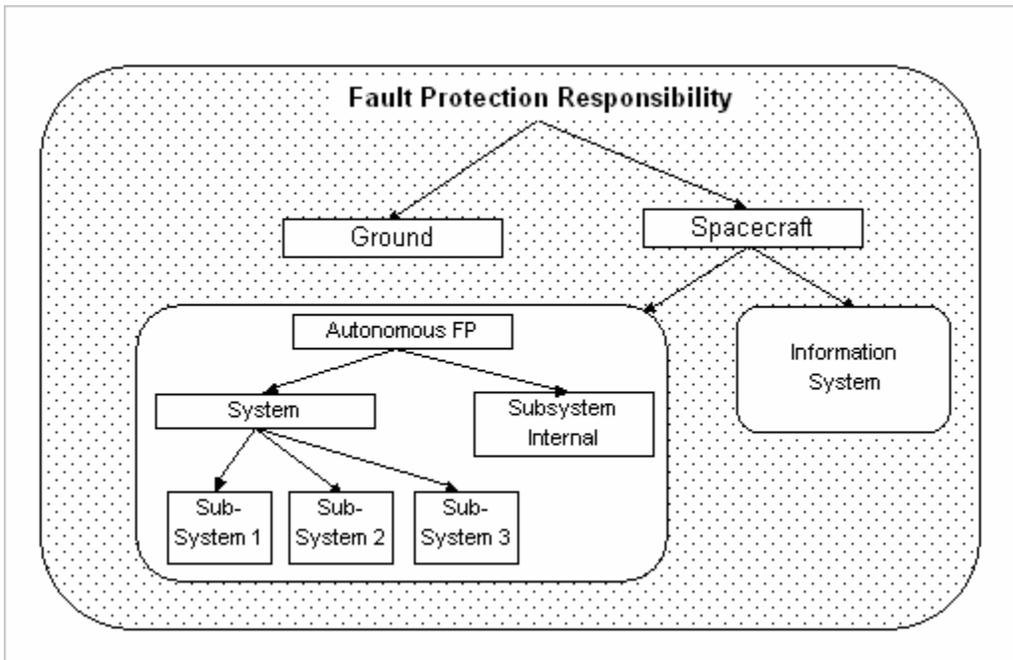


Figure 2. Fault Protection Allocation

3.1 Fault Protection Ground Rules and Requirements

In general, fault protection is designed with the following priorities:

- Protect critical spacecraft functionality
- Protect spacecraft performance and consumables
- Minimize disruptions to normal sequence operations
- Simplify ground recovery response, including provision of downlink telemetry.

On any spacecraft, fault protection ground rules and requirements are typically applied with the following principles following an anomaly: It is desirable to ensure spacecraft commandability and the maintenance of its safe state for a pre-determined period of time following any anomaly. Also, information in spacecraft telemetry that is sufficient to perform preliminary failure identification and analysis (i.e. Error Logging) is required in order for the Ground Operations Team to perform any necessary near-term actions.

The information collected by the spacecraft must be sufficient to analyze and reconstruct the sequence of fault protection events following the anomaly. Additionally, the fault protection design must be such that post-fault recovery will not require real-time ground responses in order to recover from pre-defined faults (those faults the system is designed to autonomously resolve).

In general, autonomous fault protection is usually not required to protect against spacecraft hardware and software design errors, sabotage, or operator errors, although protection against these errors is not prohibited if practical. For the design to be straightforward and fault protection actions predictable, it assumes that only one fault occurs at a time, and that a subsequent fault will occur no earlier than the response completion time for the first fault, and that multiple detections occurring within the response time are symptoms of the original fault. Also, the spacecraft hardware design ensures that a single point failure in a subsystem (including instruments) cannot propagate to its redundant unit or to another subsystem, or prevent switching to its redundant unit (i.e. fault containment).

Spacecraft faults are classified as either "interfering" or "non-interfering". A non-interfering fault is one whose fault or fault response does not compromise the integrity of the executing sequence, or the sequence does not compromise the integrity of the fault response. An interfering fault is one that does not meet the non-interfering fault definition. In general, for non-interfering faults, fault protection may execute in parallel with the executing sequence. For interfering faults, fault protection may NOT execute in parallel with the sequence. If a non-critical sequence (does not contain crucial commands) is canceled, it is not autonomously restarted. If a critical sequence (contains crucial commands) is stopped, it will be resumed after the fault response has completed. Figure 3, "Interactions with Executing (Stored) Sequences Overview" illustrates this concept.

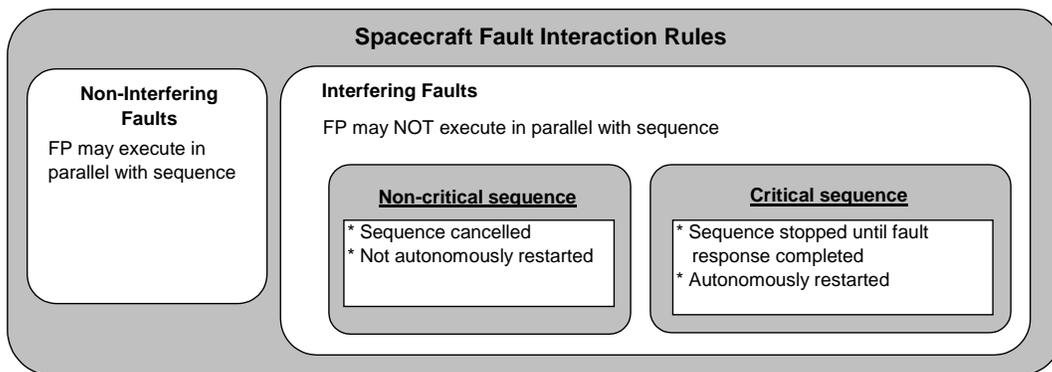


Figure 3. Interactions with Executing (Stored) Sequences Overview

If a fault is detected during critical events, fault response design must ensure the completion of the critical event as required and when required, with spacecraft safety having lower priority until the critical events are completed. Non-essential fault protection is typically disabled during critical events. To prevent the possibility of endless loop responses resulting from false alarms, fault responses are always designed within the system to be self-limiting (i.e. fault responses are eventually disabled after executing to their final actions in order to prevent continuous execution; the exception to this rule is the Command Loss Response which executes continuously until an uplink command is received).

4.0 FAULT PROTECTION ARCHITECTURE IN JPL SPACECRAFT

The main computer (referred to below as "CDS": Command and Data Processing computer), is usually the host for the spacecraft's SFP monitors and responses. SIFP is hosted within the subsystems themselves, where fault data collection/fault resolution is handled. For SFP, the CDS computer provides the required services for data collection and processing as shown in Figure 4, "CDS Services for SFP" and Figure 5 "Information Architecture for SFP" (Gavit 2003, CAS-3-331, 3-19).

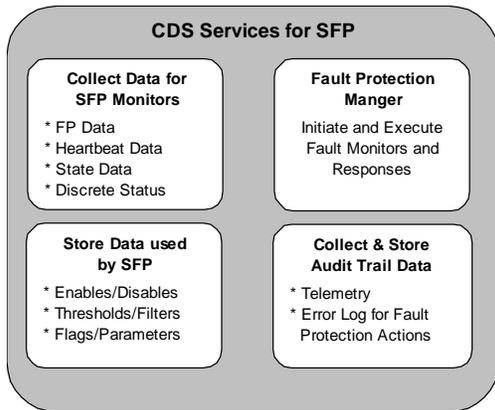
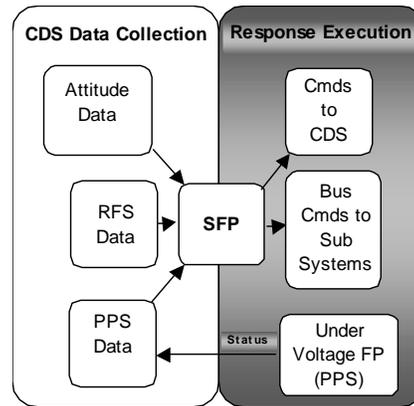


Figure 4. CDS Services for SFP



RFS - Radio Frequency Subsystem
 PPS - Power & Pyrotechniques Subsystem
 CDS - Command & Data processing Subsystem (Main Computer)

Figure 5. Information Architecture for SFP

Both SFP and SIFP are structured as a group of fault monitors and responses that are initiated and executed by their own “Fault Protection Manager” software. Monitors evaluate a spacecraft measurement against a predefined value or state to determine if a failure condition exists. If a fault condition does exist, the monitor may count consecutive occurrences to determine if the fault condition persists beyond the value of a predefined threshold. If so, the monitor will request the appropriate response.

The “predefined values” are referred to as “thresholds,” “trigger points,” or “redlines,” and represent the value at which an anomalous condition is present. The monitor design may also include logic which detects for, and ignores failed sensors. “Consecutive occurrence counters” are used in some spacecraft; these are referred to as “persistence filters” and may be used for a variety of reasons: to ensure transient occurrences do not trigger a response, to satisfy hardware turn-on constraints, or to allow other fault protection monitors to detect faults first.

Ground Operations personnel typically enable or disable monitors and responses during the mission as appropriate. This is accomplished through a software flag which may be manipulated by the Operations Team. For the most part, the fault protection is designed assuming that these flags are to be enabled throughout the mission. However, some exceptions to this strategy exist: 1) the response is only appropriate when the associated device is powered on and operating, 2) the response is required only for specific mission events, 3) the response is not appropriate for a particular event, or 4) the response is not compatible with the currently operating sequence.

4.1 Examples of Standard Fault Protection Application

4.1.1 Cassini’s Under-Voltage FP / Safe Mode FP Responses: An example of standard fault protection application is shown in Figure 6, “Standard FP Example #1: Cassini Spacecraft’s Under Voltage Fault Protection Actions for Shorted RTG” in which a Radioisotope Thermoelectric Generator (RTG) power unit (one of three on the spacecraft), has shorted out. In this example, the Power Subsystem fault protection senses a power drop below the predefined threshold for the duration of the persistence filter. The first action taken by the Power Subsystem fault protection is to diode-isolate all three RTGs, turn off (loadshed) all spacecraft non-essential loads, regain the voltage regulation, and turn on all essential hardware. It also sets three “UV Status Flags” (one for each RTG) to notify SFP that an Under Voltage event has occurred. Once the main processing computer (CDS) becomes operational, it will deliver the status of these UV Status Flags to SFP. SFP’s Under Voltage monitor will examine the state of each RTG and if enabled, will request the Under Voltage response. The response un-isolates any correctly operating RTG, unsets its UV Status Flag, and establishes a predictable, safe spacecraft state by executing the Safe Mode response.

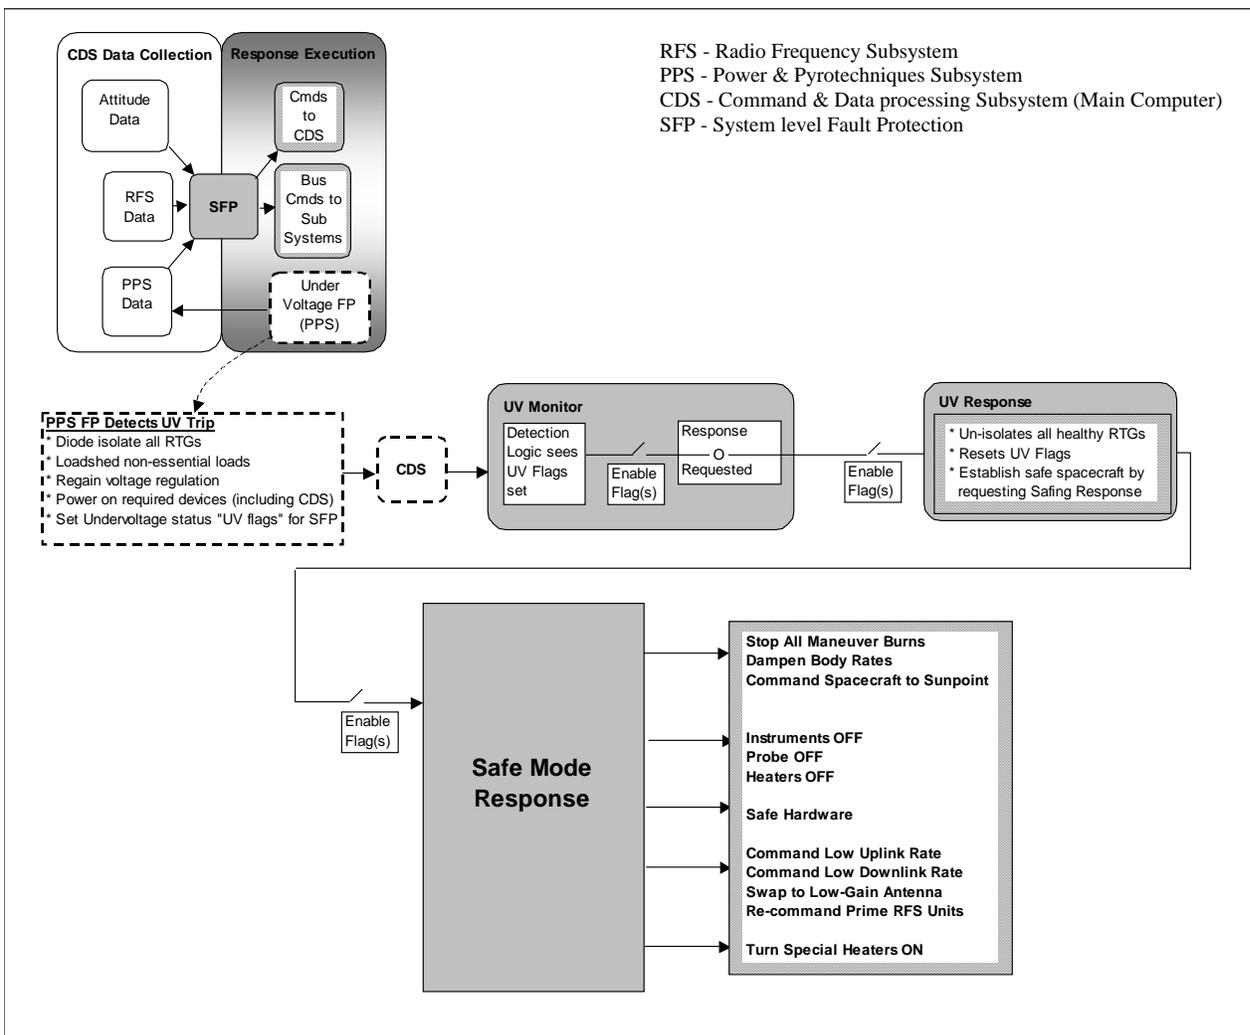


Figure 6. Standard FP Example #1: Cassini Spacecraft's Under Voltage Fault Protection Actions for Shorted RTG

4.1.2 Cassini's Command Loss Fault Response: Another example of standard fault protection application is shown in Figure 7. "Standard FP Example #2: Cassini Spacecraft's Command Loss Response". This figure illustrates how the Cassini Spacecraft's Command Loss fault protection addresses faults which inhibit communication with the spacecraft. This condition is referred to as a "loss of spacecraft commandability". For JPL spacecraft in general, the configuration of this response will depend upon the particular hardware installed. In Cassini's case, it contains dual computer (CDS) units, redundant Radio Frequency (RFS) devices, (Deep Space Transponders, Traveling Wave Tube Amplifiers, Telemetry Control Units, and three antennas (high and low gain)). The goal of this "endless-loop" response is to reconfigure the spacecraft states by performing hardware swaps and re-commanding the S/C attitude, until the uplink is restored.

A "loss of spacecraft commandability" condition is determined by a timer aboard the spacecraft which keeps track of the last time an uplink command was received from the Ground. This is a "countdown timer" that decrements continuously and is reset back to its "default value" (usually several days), each time an uplink command is received by the spacecraft. This countdown timer is the Command Loss monitor's persistence filter. The extended absence of uplink commands will eventually lead to the execution of the Command Loss Response since the timer will eventually decrement to "0". Under these

conditions, the assumption is that the spacecraft has experienced a failure and is no longer able to receive commands.

The Command Loss Response is divided up into “Command Groups” with “Command Pauses” installed after each group of commands are executed. These pauses allow several hours for the Ground Operations Team to attempt re-acquisition of the spacecraft with the newly commanded spacecraft configuration. As shown in the figure, the first Command Group will execute the Safe Mode response to turn off non-essential loads, command the spacecraft’s High Gain Antenna to the Sun, and place the spacecraft in a known uplink & downlink state (refer to Figure 6 for response actions). A 15 hour wait period is installed after this first Command Group to allow sufficient time for the Ground Operations Team to re-establish the uplink if possible. If this attempt is unsuccessful, the response will proceed with the next course of actions in Command Group #2 which will start the series of Radio Frequency Subsystem (RFS) hardware unit swaps. Five to seven hour wait periods are installed between each Command Group execution, to allow the Operations Team the opportunity to send commands to the spacecraft to re-establish the uplink with the new commanded configuration. Once the spacecraft successfully receives a command from the Ground and the uplink has been re-established, the response will halt and the “countdown timer” will be reset, leaving the spacecraft on the last commanded configuration. Since the Cassini spacecraft has two CDS main computers, a swap to the redundant unit is performed at the end of this “endless-loop” response.

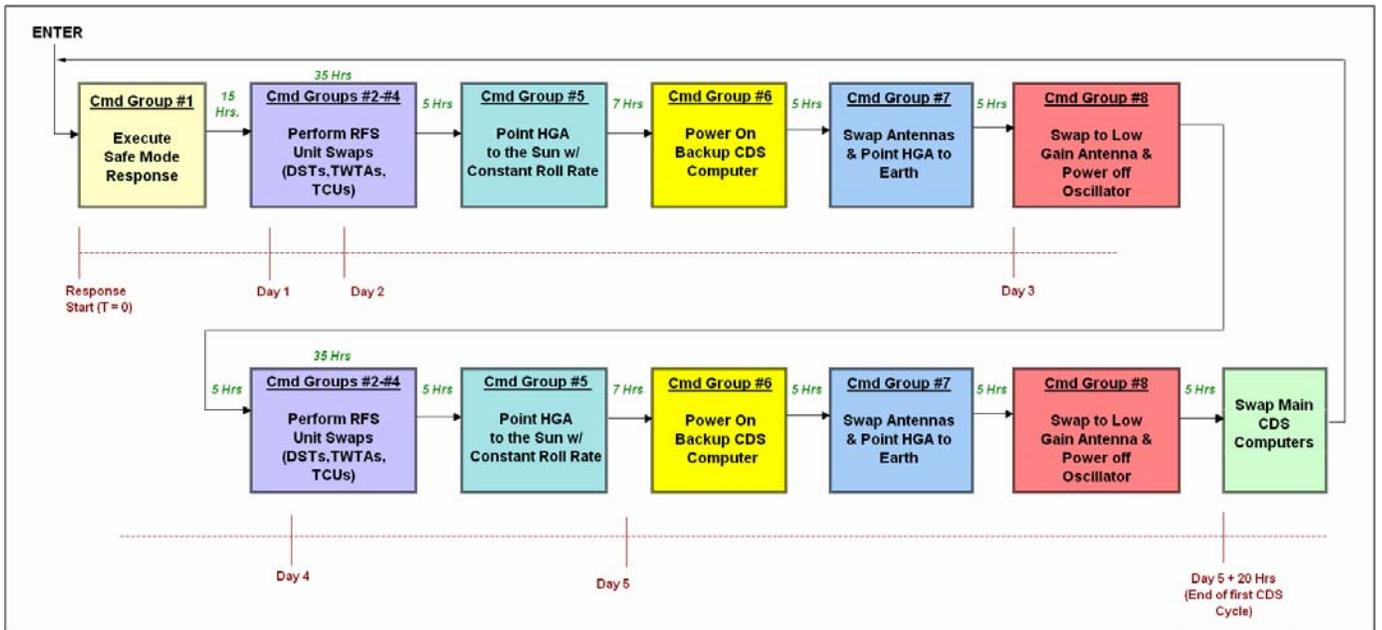


Figure 7. Standard FP Example #2: Cassini Spacecraft’s Command Loss Response

CONCLUSIONS

For spacecraft to function properly without significant risk or degradation to the mission and its objectives, continuous monitoring of the spacecraft’s components and its subsystems is desirable. An attempt to perform such a task by continuously monitoring the spacecraft’s telemetry stream is impractical, as communication through the DSN facility is quite costly, as would also be the effort of staffing people around the clock. Hence, the common problems experienced by most spacecraft: environmental influences, human error, latent component failure, fault occurrences in the presence of transmit/receive lag time, the large volume of fault possibilities due to spacecraft complexity, may be alleviated by implementing autonomous solutions within the spacecraft itself. To monitor, detect, and resolve the faults as they are encountered where possible, so that the spacecraft may preserve its overall health and provide a system with greater diagnostic capabilities.

5.0 References:

Hall, G., Schuetzle, J..1992. *Real-Time Intelligent Fault Diagnostic Systems*. Proceedings of the 1192 AIAA Space Programs and Technologies Conference

Jet Propulsion Laboratory/California Institute of Technology & National Aeronautics and Space Administration. 1993. *The NASA ASIC Guide: Assuring ASICS for Space, Appendix Three: Space Radiation Effects on Integrated Circuits*

Ong, C. Elwin. *Fault Protection in a Component-Based Spacecraft Architecture*. Massachusetts Institute of Technology. Home page. <http://sunnyday.mit.edu/papers/smcit.doc>. Accessed August 2005.

Jet Propulsion Laboratory/California Institute of Technology. 2005. *Technology section: Mars Reconnaissance Orbiter, Ulysses, Topex/Poseidon, Mars Global Surveyor, Multi-angle Imaging SpectroRadiometer, Deep Impact, Galileo, Stardust, Voyager I, Voyager II*. Home page. <http://www.jpl.nasa.gov/missions/>. Accessed August 2005.

Jet Propulsion Laboratory/California Institute of Technology. 2004. *Basics of Space Flight*. <http://www2.jpl.nasa.gov/basics/>. Accessed August 2005.

Ball Aerospace & Technologies Corp. 2001. *CLOUDSAT: Automated Fault Protection Implementation*. System Engineering Report.

Qualitative Reasoning Group Northwestern University. 2005. *Projects: DS-1 Spacecraft Principles of Operations*. Home Page. <http://www.qrg.northwestern.edu/projects/vss/docs/index.html>. Accessed August 2005.

Lockheed-Martin & Jet Propulsion Laboratory/California Institute of Technology. 1997. *STARDUST Spacecraft Fault Protection Review*.

Gavit, A. Sarah, John C. Day, Paula S. Morgan. April 14, 2003. *CAS-3-330 Fault Protection Requirements, Cassini Project, Rev C*. Jet Propulsion Laboratory/California Institute of Technology.

Gavit, A. Sarah, John C. Day, Paula S. Morgan. April 14, 2003. *CAS-3-331 Cassini Orbiter Functional Requirements Book; System Fault Protection Algorithms, Rev B*. Jet Propulsion Laboratory California Institute of Technology.

National Aeronautics & Space Administration. October 1995. *Practice No. PD-EC-1243 Fault Protection*. NASA Preferred Reliability Practices. Home Page. http://klabs.org/DEI/References/design_guidelines/design_series/1243.pdf. Accessed August 2005.