Mars Reconnaissance Orbiter In-Flight Anomalies and Lessons Learned: An Update

Todd J. Bayer
NASA Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Mail Stop 264-535
Pasadena, California, 91109
818-354-5810
todd.j.bayer@jpl.nasa.gov

Abstract—The Mars Reconnaissance Orbiter mission has as its primary objectives: advance our understanding of the current Mars climate, the processes that have formed and modified the surface of the planet and the extent to which water has played a role in surface processes; identify sites of possible aqueous activity indicating environments that may have been or are conducive to biological activity; and thus identify and characterize sites for future landed missions; and provide forward and return relay services for current and future Mars landed assets.

MRO's crucial role in the long term strategy for Mars exploration requires a high level of reliability during its 5.4 year mission. This requires an architecture which incorporates extensive redundancy and cross-strapping. Because of the distances and hence light-times involved, the spacecraft itself must be able to utilize this redundancy in responding to time-critical failures. For cases where fault protection is unable to recognize a potentially threatening condition, either due to known limitations or software flaws, intervention by ground operations is required. These aspects of MRO's design were discussed in a previous paper [Ref. 1]. This paper provides an update to the original paper, describing MRO's significant in-flight anomalies over the past year, with lessons learned for redundancy and fault protection architectures and for ground operations[1,2].

TABLE OF CONTENTS

1. INTRODUCTION

This paper picks up where my previous MRO paper [Reference 1] left off. An update on the progress of the MRO mission is provided, including some of the more notable scientific discoveries. This is followed by a description of the significant spacecraft and payload anomalies over the past year, with updates to previously reported anomalies as appropriate.

2. MISSION UPDATE

The MRO Mission has completed its Primary Science Phase of studying Mars over a full Martian year (2 Earth years), and has been approved to continue its studies for a second Mars year. The spacecraft is healthy and fully capable. MRO has far exceeded requirements for science data return: to date returning more than 70Tb of science data, compared to the requirement of 26Tb. The high quantity and quality of these data are changing our understanding of Mars in fundamental ways.

A sampling of the more notable science discoveries is provided here. References are provided to published papers, submitted papers, or abstracts as appropriate.

History of water on Mars. MRO has produced significant new evidence for widespread and long-term aqueous activity on early Mars, including discovery of extensive layered phyllosilicates underlying much of the Noachian plateau plains [Ref 2], diverse compositions of aqueous minerals scattered across Mars [Ref 3], and hydrated silicates as a third major class of aqueous minerals in addition to phyllosilicates and sulfates [Ref 4]. It has also determined that large young impact craters have apparent fluvial morphologies, indicating that ice persists in the deeper subsurface of much of Mars, even in very recent history, and providing new clues to understanding the morphologic and compositional alteration of Noachian terrains [Ref 5], [Ref 6]. MRO has helped to understand the young gully features discovered by previous missions by determining that young, light-toned gullies show no

1
[1] 978-1-4244-2622-5/09/$25.00 ©2009 IEEE
[2] IEEEAC paper #1086, Version 2, Updated December 19, 2008

evidence of aqueous minerals, and occur on slopes that are steep enough that dry flows of fine-grained materials can explain their emplacement and morphology [Ref 5], [Ref 8].

Polar Layered Terrains and Ice. For the first time, the north polar cap has been profiled in cross section by MRO's ground penetrating radar. Radar profiling has characterized subsurface layering throughout the present permanent polar ice caps and this can be related to the surface exposures as imaged by MRO's cameras. The lack of flexure beneath the polar layered deposits indicates a large present-day thermal lithospheric thickness, consistent with a chondritic abundance of internal radiogenic heat sources [Ref 9]. Packets of layers in the north polar cap are most consistent with formation during recent obliquity cycle changes, suggesting a north polar cap no younger than 10 Million years in age [Ref 9].

Present Day Climate. MRO was able observe the full course of a major planet-circling dust storm in 2007, providing a wealth of new data to understand the formation and evolution of these events and to compare with the 2001 event [Ref 11].

Northern-Southern Hemisphere Dichotomy. MRO has contributed to resolving a decades-old mystery regarding the different appearance of the northern and southern hemisphere's of Mars. Precise measurements of Mars' gravity field have revealed that the Tharsis plateau partially obscures a large elliptical basin that may be evidence for an impact origin of the hemispheric dichotomy [Ref 12]. If confirmed this structure would be the largest impact basin in the solar system.

Landing Site Certification. In support of other missions, MRO provided high resolution, 3D images to help select and certify the Phoenix landing site, and is currently doing the same for the Mars Science Laboratory, which is scheduled for launch in October 2011. MRO imaged the Phoenix lander during the parachute phase of its decent toward the Martian surface – the first time in history that a spacecraft orbiting another planet has imaged the arrival of another spacecraft. After Phoenix landed, MRO, along with the Mars Odyssey, provided regular relay service between the lander and Earth.

3. SIGNIFICANT SPACECRAFT ANOMALIES

The table below lists all the significant spacecraft anomalies since launch. Items A through I were described in [ref 1] and are not repeated here. At that time, a cause had not been determined for Anomaly H "Computer Side Swap (A->B)". This has now been resolved and is described below with new Anomaly M, "Computer Side Swap (B->A)." The other four new anomalies – J, K, L and N are described below in terms of the observables; roles played by autonomous FP and by the ground operations team; use of/impact on redundancy; root cause if known; and lessons learned.

J. HiRISE Safing after SSR full (New)

On 27 Sep 2007 spacecraft fault protection shut down the High Resolution Imaging Science Experiment (HiRISE) aboard MRO. This action was in response to the cessation of aliveness indication from HiRISE. Telemetry indicated that HiRISE stopped reporting aliveness because it had

Table 1 - Significant Spacecraft Anomalies

	Date	Anomaly	Cause
A	27-Sep-05	Memory SEUs	Redundancy scheme incompatible with unexpectedly rad-soft parts
B	2-Nov-05	Star Tracker Long Acquisition Time	Star tracker sensitivity lower than expected
C	3-Jan-06	Computer Warm Reset	FSW bug
D	26-May-06	Transponder Ka Exciter Failure	Premature part failure
E	31-May-06	Safe Mode Entry	Command sequence design error
F	26-Jul-06	Safe Mode Entry	Ground command error
G	16-Aug-06	RF Transfer Switch Failure	Most likely: RF breakdown due to flaking plating
H	14-Mar-07	Computer Side Swap (A->B)	Computer memory controller lockup
I	18-Jul-07	Payload Interface Task Suspension	FSW not fully bulletproofed to noise-induced interface data corruption.
J	27-Sep-07	HiRISE Safing with SSR Full	Defect in software/sequence interface with instrument
K	7-Nov-07	Appendage Contact With Spacecraft	Incorrect parameter combination due to incomplete requirements
L	29-Nov-07	Safe Mode Entry	Ground command error
M	14-Mar-08	Computer Side Swap (B->A)	Computer memory controller lockup
N	27-May-08	Electra UHF Anomalies	FSW bugs

internally sensed an over-temperature condition and transitioned to an internal safe mode.

Investigation determined the cause for the over-temperature condition was that commands to HiRISE had inadvertently left the instrument focal plane electronics in a high power state, causing anomalous heating. Imaging commands to instruments are contained in pre-defined sequences called 'blocks' – a strategy meant to ensure proper sequence and timing of commands. The problem was traced to an interaction between this block and the software which controls HiRISE's interface with the Solid State Recorder (SSR). The interface software contains logic to sense an SSR-full condition and to respond by preventing additional HiRISE images from being commanded until sufficient storage space becomes available. This condition was present before the anomalous shutdown. That is, the SSR full condition was detected, causing the interface software to withhold the subsequent HiRISE 'expose' command. So far this was per the design. However the design had not accounted for the existence of a certain command in the block which precedes the expose command. This command prepares the focal plane electronics for taking an image by powering the detectors but not reading them out. From a power and thermal standpoint this is identical to an imaging state. Because the interface software did not withhold this command in addition to the 'expose' command, the detectors were left in a high power state long enough to trip internal temperature limits and shut down the camera.

No damage was incurred. The block was modified to include an explicit shutdown of the detectors at the end of expected imaging, thus assuring that they are not left in a powered state. The interface with all other instruments was reviewed for similar vulnerabilities, and none was found.

Safety nets used. Fault Protection, in this case internal to the instrument, was crucial to prevent the over-temperature condition from damaging or destroying the equipment. Ground procedures were used to recover HiRISE to an operational state.

Root Cause. The interface software interactions with the command block did not provide for safely stopping an image in case the SSR was full. This off-nominal condition was not adequately tested.

Lessons Learned. Safety implications are not always obvious in complex interfaces. Therefore additional testing of off-nominal cases needs to be done in order to catch those that are not clear from analysis.

K. Solar Array Contacts Spacecraft Body (New)

The MRO spacecraft has three articulable appendages: one High Gain Antenna (HGA) and two Solar Arrays (SA). The solar arrays are designated Plus X and Minus X for the location of their mount points on the spacecraft bus. Each appendage is pointed using a two axis gimbal, each axis of which includes a motor and high resolution resolvers. The motors can be commanded to operate in a number of modes which include rate commanding, powered hold, torque commanding and disabled. For the purpose of the anomaly discussion the main mode of interest is the rate commanding mode.

The software can control the pointing of the appendages by using several articulation states. The state of relevance to the anomaly is Vector Tracking State (VTS). In this mode an appendage is commanded to track its target, either the sun in the case of the solar arrays or the earth in the case of the HGA. All of these commands must pass through the Keep Out Zone (KOZ) algorithm to keep the appendages within a desired inner/outer axis space. These allowed spaces are described by configurable line segments in the inner/outer axis phase plane, determined through analysis and verification pre-flight, and loaded onboard as configuration files.

The KOZ algorithm's function is to provide target gimbal angles to the articulation closed loop position controller which are in the allowed inner/outer space for that appendage. The KOZ algorithm does not provide a hard stop for the appendage if the inner and outer angles reach the KOZ boundary. Rather, if the measured inner and outer angles for an appendage are found to be inside the KOZ (outside the allowed area), the FSW commands the appendage to a position inside the allowed region. The algorithm performance is controlled by a set of ground-specified parameters, including maximum commanded rate and maximum commanded acceleration/deceleration.

Two-axis gimbal systems can be used to point anywhere in three-dimensional space, except when the inner gimbal axis (fixed to the spacecraft body) becomes aligned with the desired target vector. In this case the inner axis becomes useless for pointing control. This kinematic singularity is a well-known feature of two axis gimbal systems. Near the singularity, small pointing changes require that the inner axis moves a very large amount at very high speed (in the limit it must move 180 degrees at infinite speed). By design, MRO did not attempt to avoid these regions using autonomous control. Rather, the FSW controlling the gimbals simply moves the inner axis at the maximum allowed rate. This is often insufficient to maintain tracking. In the case of the HGA, this will result in loss of signal lock; therefore this situation is avoided operationally by restricting operation at attitudes and times at which this singularity occurs. In the case of the Solar Arrays, by contrast, the effect on sun tracking performance is minimal and temporary and so no operational workaround was implemented.

The Anomaly. On 7 Nov MRO performed an autonomous warm reset on C&DH B as a result of indicated failure of both redundant gimbal controllers on +X solar array. The

reset occurred during the time of a scheduled night-side SHARAD off-nadir observation, when the spacecraft downlink signal was occulted by Mars (no real-time visibility of the event). On occultation exit, the expected downlink signal was not acquired. Standard contingency plans were used to search for and find the signal at the 40bps safe mode rate. Telemetry showed the vehicle had experienced a fault protection-initiated warm reset on C&DH B after motor rate errors were reported in both primary and redundant motors. There were no indications of any damage; the boot up and safe mode configuration on side B were nominal; all three appendages had been successfully moved by fault protection to their safe mode positions; and all motors indicated 'good' condition. Recovery from safe mode was completed 14 November, and an incremental return to full science operations was completed 19 Jan 08.

The proximate cause of the anomaly is that the +X solar array violated its appendage KOZ and contacted the thermal blanket covering the spacecraft +X bay. Mechanical resistance from pushing against the blanket caused gimbal motor rate errors in the primary motor. As a result, fault protection declared the primary motor failed and autonomously swapped to the secondary. When the secondary motor encountered the same mechanical resistance, it too was marked failed. With both motors failed, fault protection initiated a heartbeat termination leading to a warm reset of the C&DH.

Damage Assessment. Following the anomaly, with the realization that the solar array had contacted the vehicle in some manner, the flight engineering team carefully examined all telemetry for indication of possible damage that may have been incurred by the event, focusing on but not limited to damage to the array, structure, blankets, and gimbals. No evidence of any damage or degradation was found.

Closer examination of historical flight telemetry revealed a total of 16 excessive incursions (defined as a KOZ incursion which exceeds 20% of the available buffer between the KOZ and the point of physical contact) starting in late summer 2007, involving one or both solar arrays, but not the HGA. All incursions occurred while the spacecraft was performing large science rolls, including both daytime and nighttime rolls. Each of these incidents was analyzed with the spacecraft structural modeling program to quantify how close to the spacecraft bus or HGA the solar arrays had come. This analysis confirmed that none of the previous incursions resulted in a contact event.

An analysis was performed of the orbit geometry and event timing associated with each of the incursions. Two factors were found to be associated with all of the excessive incursions: the incursions all occurred at a point in the mission near the minimum sun-beta angle and all of the off nadir slews were performed at a time and in a direction that put a singularity axis of a solar array close to the sun vector.

Causal Chain. Investigation determined that the solar array violated its KOZ because the combination of spacecraft-Mars-Earth geometry and slew timing created conditions where the appendage was moving near a kinematic singularity, causing the appendage to encounter the KOZ boundary while moving at its maximum allowed rate. When travelling at this rate, the maximum allowed deceleration was insufficient to arrest the appendage motion before it contacted the spacecraft. Put simply, the parameter values for the combination of maximum rate and maximum acceleration were incorrect.

These parameters were loaded prior to launch and never modified in flight. The parameters were selected to satisfy appendage tracking and spacecraft pointing stability requirements. The effect of this combination of parameters on KOZ enforcement was considered in the original version of the parameters, but later, when the final pre-launch change was made, the effect was missed.

Also, the effect of these parameters on KOZ penetration escaped detection in the pre-launch verification and validation process.

Root Cause. No single root cause, by itself, was found to completely explain the anomaly. The investigation identified several root causes, any one of which, if fixed, could have prevented the problem:

(1) The spacecraft system level requirements for configurable appendage KOZs were incomplete with respect to the elastic nature of the KOZ.

(2) The documentation in the parameter database regarding the critical interaction of these two parameters was insufficient to alert engineers to the risk of changing one without changing the other.

(3) Validation activities failed to catch the fact that requirements and verification were incomplete.

Contributing Factors. Several factors were identified which contributed to the anomaly

(1) Misleading terminology contributed to miscommunication between system and subsystem engineers. Pervasive and consistent use of the term "Keep Out Zone" did not adequately convey the elastic nature of the boundary.

(2) Incomplete system-level evolution of gimbal control concepts from simpler heritage systems contributed to an incomplete system KOZ requirements specification. The change from hard and soft stop implementations to a design featuring two-axis software-controlled

motion with vector tracking while slewing was not fully internalized.

(3) Acceleration and rate limit parameters were used for multiple and conflicting purposes (to minimize perturbations; to maximize tracking performance; and to constrain motion).

(4) Original conservative off-nadir roll analysis was deemed too conservative and modified, but the implications of the original assumptions were not preserved or carried into operations.

(5) Incomplete knowledge transfer between individuals and teams inhibited universal understanding of the KOZ implementation.

(6) The in-flight telemetry for KOZ performance was insufficient to indicate depth of KOZ penetration.

Corrective Actions for MRO. Several changes were made to prevent a recurrence. First the gimbal appendage rate parameters were modified on-board to limit the travel into the KOZ under worse-case conditions. Second, the ground verification process was changed to include a complete ground simulation run prior to uplink of each week's science activities. Third, additional telemetry points were defined for enhanced ground monitoring capability. Additionally, a flight software patch was considered – but ultimately judged unnecessary – which would have provided an additional backstop against KOZ violations.

Safety Nets Used. Spacecraft Fault Protection was critical to preventing damage by detecting and responding to gimbal motion problems. Although redundant motor controllers were used, this redundancy played no significant role in surviving the fault. Finally, ground procedures were necessary to find the spacecraft safe mode downlink and recover from safe mode.

Lessons learned. Eight significant Lessons Learned resulted from this anomaly.

(1) Requirements terminology must be consistent with the design implementation.

(2) Future systems should consider requiring a hardware or software "hard stop" that positively prevents collisions for moving parts.

(3) When approaching complex problems using analysis, use of simplifying conservative assumptions is usually the best approach, even if this means not being able to maximize system performance. Until less conservative assumptions are verified by higher fidelity analysis, the implications of reduced conservatism need to be carried as risks.

(4) Misleading terminology is always a problem. Better terminology could have led to a different implementation (SW hard stop) or to a requirement to quantify the maximum excursion into the KOZ. Precision in requirements language is important, and important to correct even though late in development it may seem like "wordsmithing."

(5) MRO's Parameter Database is a powerful ground tool allowing the MRO team to track many thousands of parameters, allowing the ability to designate certain parameters as mission critical, and the ability (and requirement) to include key notes for all parameters. The gimbal rate and acceleration limits should have been designated as such and documented accordingly. Discipline in completing such databases is difficult but critical, especially when there are thousands of parameters to describe and review. In addition to the subsystem Cognizant Engineers and Flight Software engineers, a knowledgeable System Engineer is a key participant in this process.

(6) Having complete telemetry to fully monitor, trend and alert operations personnel of out-of-bounds performance is essential for critical and complex functions.

(7) The institution should consider incorporating in the JPL Design Principles these findings regarding design and verification of configurable appendage articulation keep out zones.

L. Command Error Induces Spacecraft Safing (New)

On the afternoon of 11/29/2007, MRO entered safe mode. The cause of the safe mode entry was an uplink error in which a command file to power on an instrument was inadvertently sent twice. The power on sequence calls an on-board block to power on the instrument and this block was already active when the second uplink reached the spacecraft, resulting in an execution contention whose fault protection response was to command a safe mode entry. The vehicle entered safe mode nominally and ground operations personnel recovered the spacecraft to nominal mode the same day. Resumption of science operations occurred later, after process improvements were implemented.

In response to this command error, the project reviewed its procedures for spacecraft commanding and made several changes to the procedures and protocols used by the operations engineer known as the "Ace". The Ace is the person who actually transmits commands to the spacecraft. These changes were in two areas: reducing the number of parallel activities the Ace must manage and coordinate, and mandating additional checks at key points in the commanding process. As longer term fixes, the ground software tools used by the Ace during commanding were reviewed and strengthened with additional error prevention features.

Safety Nets Used. Onboard fault protection worked as designed by preventing an indeterminate state which could have been caused by improper commanding. Ground personnel and procedures were necessary to recover the spacecraft to nominal operations.

Root Cause. Human error, contributed to by excessive demands on attention, and inadequate safeguards.

Lessons Learned. Positive lesson is that fault protection was correctly designed to prevent damage due to human error. Negative lesson is that the error was preventable by better safeguards and more reasonable tasking.

M. Spacecraft Computer Side Swap (B->A) (New)

On 14 March 2007, MRO performed an unrequested warm reset followed by an unrequested side swap to C&DH Side B. The investigation at that time concluded that insufficient information was available to determine proximate cause, or to rule out a permanent failure of Side A. The investigation was closed in September 2007 with eight most-likely proximate causes. See [Ref 1]. A deliberate decision was made to leave enabled the fault protection autonomous side swap capability. This was for three reasons: 1) even if side A were permanently failed, no failure mode postulated would prevent the autonomous return to side B; 2) Side A was thought likely to be fully functional; and 3) a side swap is the only way in the MRO architecture to effect a cold reset of the C&DH.

On 13 February 2008, MRO performed another unrequested warm reset followed by an unrequested side swap – this time back to C&DH Side A. This is referred to as Side Swap #2. The investigation was reopened and benefitted from the significant new information. First, we could immediately rule out permanent hardware failure as the cause of Swap #1 since Side A was functioning correctly. Second, whatever occurred to cause the first swap was cleared by a power cycle of the Side A computer. Third, the observables from both swaps were essentially identical, strongly suggesting they had the same cause. This gave us confidence that Side B would also be functional were we to swap back to it. This also allowed us to further constrain potential causes to those that could occur on two independent sets of hardware.

With this information, the existing fishbone was reviewed again. Five of the eight most likely proximate causes from 1st swap were eliminated, being permanent hardware failures. Only software flaws remained, but all still seemed unlikely.

From the previous investigation it was known that a sufficiently large number of errors in memory would delay boot-up long enough to trigger a heartbeat timeout and cause a side swap. But no plausible mechanism to create these errors was known. A key insight occurred during the early Swap #2 investigation: if the background process which refreshes the dynamic RAM memory were shut off prior to the boot, this would cause enough errors and provide the missing mechanism. Given this insight, all potential faults were re-analyzed. In particular, we re-reviewed the vendor-published problem list associated with BAE's cPCI RAD750 spaceflight computer (referred to as the RAD750 Errata). In reviewing this problem list, we found one which would indeed have the side effect of halting memory refresh. This became the leading candidate. The main analysis/test efforts were focused on creating this fault in the OTB, and other efforts worked toward ruling out the other potential causes. The investigation has now concluded that this known problem was the most likely proximate cause.

Root Cause. The flaw was discovered by another BAE customer after MRO had launched. MRO examined this new flaw in 2006 when it was published and determined that no action was necessary. This was due to the erroneous understanding that occurrence of this unlikely event would only cause a warm reset and not a side swap. This was because MRO had an insufficient understanding of the low level details of the operation of the CPU under these circumstances and how that impacts the MRO system design. It is not clear from the Erratum that memory refresh is terminated during the halt. The BAE description relies upon the RAD750-internal Watchdog timeout to recover the system as this mechanism is common to all RAD750 customers, and the recovery after this watchdog timeout restores refresh. MRO implementation set the RAD750 watchdog longer than the spacecraft Fault Protection heartbeat watchdog timeout in order to allow system fault protection mechanisms to manage reboots. The combined effect of memory refresh termination and the system watchdog timeout period was not considered. Without refresh, many thousands of memory errors accumulated before the timeout expiration caused a warm reset. MRO startup was not designed to handle this many memory error during a warm reset. The MRO approach to warm boot attempts to recover memory contents, which requires a non-destructive memory test. Non-destructive Memory Test takes too long handling memory errors, resulting in a 2nd timeout and a side swap. MRO analysis of Erratum 24 missed this subtlety. Note that a cold reset does not try to recover memory contents, initializes memory before testing, and therefore works.

Corrective Action. A known workaround exists for the published problem, which MRO has now implemented. No further side swaps are expected due to this issue. Other projects using the BAE cPCI RAD750 architecture should implement the workaround as well.

Safety nets used. As in the first swap, it was the low level firmware-based FP which was critical to detection of a failure to boot and forcing a swap to the backup computer. High Level FP Software was then necessary to configure

the vehicle in safe mode. Ground contingency plans for LOS were necessary to find the safe mode signal, and Safe Mode Recovery Contingency Plan was necessary to return the vehicle to nominal operations. After the first side swap, redundancy was believed to have been critical to saving the mission. However now that proximate cause is known it is clear that the crucial factor was an autonomous cold reset. Redundancy of C&DH for this fault is only required because in the MRO architecture swapping sides is the only way to effect a cold reset.

Lesson learned: Several significant lessons learned resulted from the side swap anomalies:

(1) Future designs should increase the amount of low level information available to diagnose boot-up problems, for example enhancing the amount of boot trace data available.

(2) Ensure all areas of all memories are robustly error-checked. This is not related to proximate cause, but the investigation revealed some memory which was not checked, and some checks on other memory which were not fully robust.

(3) Carefully examine trade between recovering data on warm reset vs. complete initialization. Data salvage requires more work to be robust against the type of memory errors encountered in this anomaly.

(4) Missions need a clear-everything capability to completely reset all FPGA and ASIC logic. MRO's side swap provides this capability, which proved crucial to recovery.

N. Electra UHF Relay Anomalies (New)

Overview of a Relay Overflight. A nominal relay overflight between MRO and Phoenix lasts approximately 35 minutes in terms of MRO activities. The actual period where the two craft are in contact is shorter: up to about 15 minutes. The orbiter, in this case MRO, initiates the relay session by commanding the Electra UHF Transceiver (EUT) to begin sending 'hail' signals to the lander. When the lander receives this signal it responds and using a standard protocol the orbiter and lander 'handshake' to establish the agreed link configuration. This is followed by exchange of data: the orbiter sends commands to the lander which it had previously received from Earth ("forward link"); and the lander sends engineering and science data to the orbiter for return to Earth ("return link"). The length of each overflight is agreed in advance by MRO and Phoenix controllers and uplinked as command sequences to MRO. At the commanded end-of-session, MRO commands the EUT to close the link, and Phoenix shuts its link down after detecting the end of MRO transmissions.

The EUT anomalies began on 27 May 2008, the second Martian day (Sol) after Phoenix landing. From then through 11 August, a total of 163 overflights have been executed. Of these, a total of 8 have been anomalous, and 5 of these anomalous passes have resulted in PHX data loss. The set of anomalies exhibited four distinct signatures. These anomalies were the subject of an intensive investigation within the MRO Project and the Mars Exploration Directorate.

Type 1 ("Heartbeat Anomaly"). The first type has had a single occurrence – the first anomalous overflight. Electra heartbeat timestamp stopped updating during pre-session setup. EUT seemed to continue processing commands until S/C powered Electra off per expected response to time not updating.

Immediate Operational Responses. After this first anomaly, the Flight Team powered on and reconfigured the EUT in time to support the next overflight, while Electra team began analyzing data from the anomaly. During that next overflight, the EUT experienced its second anomaly, which was later classified as the first of the "Type 2" anomalies (see below). Meanwhile, Phoenix switched to Odyssey for all relay support until MRO's anomalies could be resolved.

Type 2 ("Safe Mode Anomaly"). Four occurrences. Electra unexpectedly went into safe mode during repeated hailing. EUT heartbeat stopped updating prior to this, and the reboot appeared to produce no core dump.

Type 3 ("Out-of-frame-sync Anomaly"). Three occurrences. Data exchange stopped shortly after communication established. Electra continued counting "out of sequence" frames. This has only occurred at start of some passes at low elevations. This was found to have the same signature as an anomaly seen on the EUT Engineering Model in the MRO Orbiter Testbed during Operational Readiness Test #6 in early 2008. No root cause was found at that time; it was concluded to be most likely a test set up problem, or at worst a rare fault.

Type 4 ("End of pass anomalies after reboot"). Three occurrences, associated with the first three Type 2 anomalies (fixed prior to the fourth Type 2 occurrence). Following a type 2 unexpected reboot anomaly, spacecraft commands do not execute or execute anomalously, accompanied by strange telemetry readings. This only occurs after Electra reboots while the S/C continues activities unknowingly.

The team evaluated the full gamut of proximate causes, including: sequence/command error, hardware, software, hardware/software interactions, environmental, interface faults, RF link itself, clocks, and others. After determining proximate cause, we then performed a root/contributing cause analysis for each anomaly type. We analyzed a fifth anomaly type, which we called "Aggregate Anomaly", to help us understand how four separate defects could have escaped undetected. The following summarizes the

proximate and root causes (contributing factors are omitted for brevity) and lessons learned.

Type 1

Proximate Cause: EUT flight software (FSW) bug causes task to hang.

Root Causes: insufficient breadth/depth in software peer reviews; lack of semaphore analysis for this particular semaphore; independent verification/ validation (IV&V) did not check for condition; institutional design principles lacked specific guidance; incorrect initial EUT memory sizing requirements led to triggering bug in flight.

Type 2

Proximate Cause: Undocumented behavior of microcontroller which causes it to hang up in certain nested interrupts.

Root Cause: unclear processor documentation regarding functionality of nested interrupts.

Type 3

Proximate Cause: Protocol implementation flaw in the lander radio.

Root Causes: incomplete protocol implementation on Phoenix side of UHF link; insufficient documentation of MRO – Phoenix compatibility issues; lack of a standard design practice for command parsing; insufficient box level testing under realistic conditions including marginal links; insufficient rigor in designing comprehensive system test program; insufficient resources to investigate a prior occurrence of the anomaly during ground test.

Type 4

Proximate Cause: EUT FSW bug causes it to lose synchronization with the spacecraft command interface.

Root Causes: inadequate vendor documentation of EUT command/data interface chip; inadequate loop-closure in MRO to EUT interface design; insufficient documentation of test requirements on MRO relay pass on-board command sequence.

Aggregate Anomaly

No single root cause was identified which would have prevented all four of the anomaly types. However several contributing factors were identified and helped form the set of lessons learned.

Operational workarounds as well as fixes to flight software were implemented for each of these anomaly types. MRO was able to resume support to Phoenix with no further anomalies, and supported through the end of the Phoenix mission in November 2008.

Safety Nets Used. In the Type 1 anomaly, S/C Fault Protection was necessary in order to recognize and respond to a non-responsive EUT, and ground procedures were necessary to recover the EUT to normal operations. The other anomaly types did not entail fault protection action or ground procedures. Also, since in no case was there a hardware problem, use of the redundant EUT was unnecessary.

Lessons Learned. After the team completed the root cause analysis, they developed a set of specific process improvement recommendations. These are organized into broad categories and summarized below.

Software Design Reviews and Analysis

- Improve rigor of software peer review process
- Perform inheritance reviews of interface-related hardware and software
- Re-evaluate software IV&V process

Interface Specification

- Improve interface control document rigor regarding protocol compatibility
- Tighten EUT-to-spacecraft control loops in future
- Use fault trees analysis for critical protocols function verification

Testing and Testbeds

- Increase systems/subsystems interaction in creating system level Verification/Validation plan
- Involve domain experts earlier in low level interface testing
- Box/subsystem test liens must be tracked and carried into systems program
- On-board relay sequences/blocks require comprehensive verification matrix
- Apply greater emphasis on off-nominal/stress testing early, during unit testing
- Investment is needed to increase the fidelity and usability of box-level testbeds

Management

- Given that EUT is a critical infrastructure element for the Mars Exploration Program, a better management model for EUT development, in terms of oversight and support, would be a Command & Data Handling subsystem rather than a science payload.

Resources

- Sustaining engineering should be provided on any critical subsystem, especially one with

evolvable and reconfigurable characteristics intended to support multiple missions.
- Apply greater resources during development to prevent Electromagnetic Interference problems

State of the Practice (Principles, Methodologies, Tools)

- Formulate institutional design guidance on semaphore management, robust command parsing, and implementation of widely used interface protocols.
- Improve process for evaluating/certifying vendor provided parts and software
- Enable institution-wide information sharing on problems and idiosyncrasies with widely used parts, units and software (e.g., a Wiki)
- Improve institutional problem reporting tools to improve cross-project collaboration
- Increase use of system modeling and use cases during requirements development in order to help validate requirements and to drive out off-nominal test cases
- Devise methodologies to help projects cope with concurrent development of highly-interfaced components and spacecraft.

In addition to the above, the investigation team developed specific implications and recommendations for other users of Electra technology and conveyed them directly to those projects.

4. SIGNIFICANT PAYLOAD ANOMALIES

The table below lists all the significant payload anomalies since launch. Items A through F were described in [Ref 1] and are not repeated here. At that time, Anomaly D (Unexpected EMI on Electra) was not fully closed; a cause had not been determined for Anomaly E (MCS Position Errors); and Anomaly F (HiRISE Detector Degradation) was understood but was being watched and trended by the project. A brief update on these three is provided below, followed by an examination of one new payload anomaly, #G (CRISM Cryocooler Anomalies).

D. EMI on the Electra UHF Radio (Update)

The unexpected EMI on the Electra UHF radio was believed to have been mitigated sufficiently to allow MRO support to the Phoenix mission. With the successful support during Phoenix's prime mission, this belief has been confirmed. Although other anomalies occurred with the radio – described in this paper – the interference issue was proven to be resolved.

E. MCS Position Errors (Update)

Intermittent telescope actuator position errors have continued. No root cause has yet been determined, but the evidence points to mechanical contamination of some kind. The working hypothesis is that a reservoir of contamination exists in a region of the mechanical gear train, and that this can introduce new particles over time which get transported throughout the range of motion of the mechanism, causing new position errors at new locations, which eventually disappear in a way consistent with them being 'ground up'. The MCS team has operated in a tactical mode to tweak instrument scan patterns following each new error in an attempt to avoid the new error location. While this is arduous -- and a distraction from the planned scientific analysis – the approach has allowed recovery of most of the scientific goals of the investigation.

F. HiRISE Detector Degradation. (Update)

The rate of degradation has been consistent with predictions: very low and manageable. Overall, HiRISE has continued to perform in an outstanding fashion. Meanwhile the root cause of the ADC impedance drift has been determined to be trace contamination of the part during manufacture. This is not considered a serious issue because it does not impair the part's function, but only causes a change in certain characteristics, and these characteristics are only important because of an unrelated oversight in the original circuit design of the camera focal plane. Indeed, ground testing has shown that operating at warm temperatures can redistribute the Chlorine contaminant and actually reverse the degradation. This mitigation is being

Table 2 - Significant Payload Anomalies

	Date	Anomaly	Cause
A	3-Nov-05	HiRISE Sunshade Low Temperature	Sunshade blanket design and outdated thermal model
B	13-Dec-05	MCS Anomalous Power-up	Inadequate test fidelity
C	28-Sep-06	SHARAD Safing	Command sequence error/inadequate instrument software response.
D	7-Nov-06	Unexpected EMI on Electra	EMI from gimbaled payload had undiscovered position dependence
E	11-Dec-06	MCS Position Errors	Most likely: debris in mechanism
F	17-Jan-07	HiRISE Detector Degradation	Insufficient design margin / part contamination
G	29-Apr-08	CRISM Cryocooler Anomalies	Currently unknown

implemented onboard in order to increase HiRISE's useful life.

G.CRISM Cryocooler Anomalies (New)

The Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) provides high resolution imaging (18m) over 544 spectral bands, for unprecedented observations of Martian mineralogy. A Stirling-cycle cryogenic cooler provides the low temperatures required to observe in near-infrared bands. The cooler has a limited lifetime; therefore CRISM carries 3 of them. Late in 2007, Cooler #1 was shutdown autonomously by the instrument after exceeding its current limit. An investigation concluded that the cooler was operating in a hotter than predicted environment, causing the unit to work harder than expected, and that this probably caused a premature failure. The operating conditions for the remaining two coolers were revised to preclude additional failures. This involved using a higher temperature setpoint to reduce the load on the coolers and earlier and longer de-icing cycles, both at some cost to science quantity and quality.

Later, coolers 2 and 3 exhibited different anomalous behavior. Each cooler at separate times was unable to fully achieve the temperature setpoint. Initially, loss of He (the coolant fluid) was suspected. But this was discounted when later performance returned to normal.

The investigation of Cooler 2 and 3 behaviors is ongoing, with the two leading theories being:

(1) non-uniform lubrication that improves when the cooler runs at a lower setting and

(2) contaminant(s) in He that collect/gets-dispersed with time. In the meantime a number of cooler management strategies have been implemented/proposed to preserve longevity.

Safety Nets Used. The existence of redundancy in the life-limited cryo-coolers was crucial to preserving CRISM science capability. Ground procedures were developed to work around the anomalies and manage the remaining cooler lifetimes. In the case of the cooler 1 failure, internal fault protection protected the instrument from being damaged due to over-current.

Root Cause and Lessons Learned. Not known at this time.

5. CONCLUSION

MRO has been a highly successful mission and has every expectation of a long and productive extended mission. Problems have continued to crop up in the second year of the science mission. In surviving these anomalies, MRO continues to prove the value of a well-crafted safety net woven from the strong threads of redundant hardware, solid autonomous fault protection, and a prepared and alert ground operations team.

The lessons learned from our anomalies are helping MRO to avoid similar problems in the future. It is hoped that other missions in development will also find them valuable. See references [13] and [14] for more information and updates on the MRO mission.

ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The author would like to acknowledge the critical and outstanding contributions of the MRO spacecraft Flight Engineering Team at Lockheed Martin Space Systems Company, and the very helpful review of this paper's Mission Update section by Dr. Richard Zurek, MRO Project Scientist.

REFERENCES

[1] Bayer, T. J., "In-Flight Anomalies and Lessons Learned from the Mars Reconnaissance Orbiter Mission," IEEEAC Paper #1451, 2007.

[2] Mustard, J., et al. (2008), Hydrated Silicate Minerals on Mars Observed by the CRISM Instrument on MRO, Nature, 454, 305-309,10.1038/nature07097.

[3] Ehlmann, B.L., J.F. Mustard, J.L. Bishop, G.A. Swayze, L.H. Roach, R.N. Clark, R.E. Milliken, F. Poulet, S.L. Murchie, and the MRO CRISM Team (2008), Distinct provinces of aqueous alteration in the western Isidis region identified with MRO-CRISM, Lunar Planet. Sci 39, 2326.

[4] Milliken, R. E., G. Swayze, R. Arvidson, J. Bishop, R. Clark, B. Ehlmann, R. Green, J. Grotzinger, R. Morris, S. Murchie, J. Mustard, and C. Weitz (2008), Opaline Silica in Young Deposits on Mars, Geology, in press.

[5] McEwen, A.S., et al., (2007), A Closer Look at Water-Related Geologic Activity on Mars: Science, 317, 1706-1709, doi: 10.1126/science.1143987.

[6] Mouginis-Mark, P.J. and H. Garbeil (2007), Crater geometry and ejecta thickness of the Martian impact crater Tooting. Meteoritics and Planetary Science, 42, 1615 - 1626.

[8] Pelletier, J.D., K.J Kolb. , A.S. McEwen, and R.L. Kirk (2008), Recent bright gully deposits on Mars: Wet or dry flow? Geology, 36; no. 3; p. 195-198; doi: 10.1130/G2 4346A.1

[9] Phillips, R. J., et al. (2008), Mars North Polar Deposits: Stratigraphy, Age, and Geodynamical Response. Science, 320, pp. 1,182-1,185.

[11] Cantor, B.A., and Malin, M.C. (2007), Martian weather: Approximately 5 Mars years of MOC and MARCI observations, American Astronomical Society, DPS Meeting 39, Abstract No. 17.01, Orlando, Florida.

[12] Andrews-Hanna, J. C., M. T. Zuber, and W. B. Banerdt (2008), The Borealis basin and the origin of the martian crustal dichotomy, Nature, 453, 1212-1215, doi:10.1038/nature07011.

[13] Zurek, R. W. and S. E. Smrekar (2007): An overview of the Mars Reconnaissance Orbiter (MRO) science mission. J. Geophys. Res., 112, E05S01, doi:10.1029 /2006JE002701.

[14] NASA's MRO website: http://mars.jpl.nasa.gov/mro/

BIOGRAPHY

Todd J. Bayer *received his B.S. in Physics in 1984 from the Massachusetts Institute of Technology. He started his career as a project officer in the US Air Force at Space Division in El Segundo, California. Following his military service, he joined the staff of JPL in 1989. He has participated in the development and operations of several missions, including Mars Observer, Cassini, and Deep Space 1. During a leave of absence from JPL, he worked as a systems engineer on the European next generation weather satellite at EUMETSAT in Darmstadt, Germany. He was the Lead Flight System Engineer for MRO's development, and at launch he assumed the role of MRO Chief Engineer which he held until October 2008. Currently he is the Assistant Manager for Flight Projects of JPL's Systems and Software Division.*