Soil Moisture Active Passive Mission: Fault Management Design Analyses

Peter Meakin[1], Raquel Weitl[2]
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California

As a general trend, the complexities of modern spacecraft are increasing to include more ambitious mission goals with tighter timing requirements and on-board autonomy. As a byproduct, the protective features that monitor the performance of these systems have also increased in scope and complexity. Given cost and schedule pressures, there is an increasing emphasis on understanding the behavior of the system at design time. Formal test-driven verification and validation (V&V) is rarely able to test the significant combinatorics of states, and often finds problems late in the development cycle forcing design changes that can be costly. This paper describes the approach the SMAP Fault Protection team has taken to address some of the above-mentioned issues.

Nomenclature

FM	=	Fault Management
FMECA	=	Failure Modes Effects and Criticality Analysis
FTA	=	Fault Tree Analysis
FP	=	Fault Protection
FPGA	=	Field Programmable Gate Array
FSW	=	Flight Software
GNC	=	Guidance, Navigation, and Control
IRU	=	Inertial Reference Unit
LEO	=	Low Earth Orbit
LV	=	Latch Valve
MM	=	Mitigation Matrix
NEN	=	Near Earth Network
OBS	=	Observatory
PRT	=	Platinum Resistence Thermometer
RAD	=	Radiometer
RCS	=	Reaction Control System
RF	=	Radio Frequency
RPM	=	Revolutions per minute
RWA	=	Reaction Wheel Assembly
SAR	=	Synthetic Aperture Radar
SMAP	=	Soil Moisture Active Passive
TTC	=	Time to Criticality
V&V	=	Verification and Validation
WU	=	Wheel Unit

I. Introduction

The Soil Moisture Active Passive (SMAP) spacecraft studies the Earth soil moisture and freeze / thaw state via an active L-band 1.26 Ghz Synthetic Aperture Radar (SAR) and a passive L-band 1.4 GHz Radiometer (RAD) instrument. Over its three year mission, the SMAP mission completes a global map of the Earth surface every two to three days with an eight day repeat ground track. In order to achieve full land coverage with the SAR and RAD

[1] SMAP Fault Protection Lead, Fault Protection and Autonomy Group, 4800 Oak Grove Dr., M.S. 321-450G, Pasadena, CA 91109-8099, Peter. C. Meakin@jpl.nasa.gov
[2] Fault Protection Engineer, Fault Protection and Autonomy Group, 4800 Oak Grove Dr. Pasadena, CA, raquel.m.weitl@jpl.nasa.gov

instruments, SMAP employs a rotating six-meter mesh reflector with a rotation rate of approximately 14.6 RPM creating wide swaths of coverage[1]. SMAP is a three-axis stabilized spacecraft in a sun-synchronous low earth orbit (LEO) at 685 km altitude. The orbit is a 6am 6pm ascending-descending node that experiences seasonal eclipses on the order of 20 minutes. The significant amount of stored momentum in the spinning reflector is balanced by four reaction wheel assemblies (RWAs) which are used for both momentum compensation as well as fine attitude control of the near zero momentum spacecraft. A reaction control system (RCS) of eight hydrazine thrusters is used to perform delta-V maneuvers for orbit maintenance and can be used to provide three-axis attitude control. Spacecraft power generation is provided via body fixed solar arrays and a secondary battery provides energy storage during periods of eclipse.

SMAP is a class C mission with largely single string hardware and some limited hardware redundancy. The fault management (FM) strategy is to employ a limited on-board fault protection (FP) system which is required to tolerate a protected fault set. The flight software (FSW) based FP design detects errors and takes corrective actions that recover enough functionality to place the observatory into a safe state and await ground instructions. Although the SMAP orbit provides frequent ground contact opportunities with the near earth network (NEN), attitude control is required to point the spacecraft into an orientation which supports a predictable RF communication link and simultaneously provides a power positive attitude with the solar arrays pointed toward the sun. Of particular significance to the on-board FP design is the large momentum stored in the spinning antenna and RWAs and the tolerance to failure scenarios that result in the spin-down of the antenna. Given the relative complexity required to recover active attitude control, and the limited staffing available on a class C mission, the FP team endeavored to perform as much early design analysis and validation as possible in an attempt to streamline the implementation and verification phases of the mission.

II. Fault Management Design Challenges

Fault Management, as a discipline, encounters several classical design challenges that are often not addressed until the formal V&V campaign. The first observed deficiency is that although system-level fault analyses such as Failure Modes Effects and Criticality Analysis (FMECA) and Fault Tree Analyses (FTA) are performed, there are rarely explicit, verifiable connections between the fault analyses and design requirements on the flight and ground fault management design. Of particular weakness is the connection and flow-down of fault management centric requirements on the mission system requirements on the ground operators. Secondly, after a “protected fault set” has been generated, it is rare that a rigorous analysis is performed to ensure that the flight (and ground) fault management design detects the resulting errors and responds with the appropriate mitigation before the effects of the fault result in a critical failure effect (such as unacceptably low battery state of charge). Engineering judgment is often applied during design specification to identify timing scenarios, which are deemed “stressful”, but rarely is a systematic approach applied for the entire protected fault set. Lastly, the interaction between autonomous FP responses, ground interactions, and other software behaviors is rarely evaluated explicitly. Typical practice has been to find these “idiosyncrasies” during the test campaign rather than model and evaluate them at design time.

III. SMAP Fault Management Design Approach

In order to understand the failure modes and the effects on the system, SMAP performed a Functional Failure Modes, Effects and Criticality Analysis (FMECA) and a Mission level Fault Tree Analysis (FTA). These analyses are typical reliability products developed during early mission phases for JPL missions. The FTA is a “top down” analysis, which focuses on the failure modes of the high-level mission phases and subsystem behaviors. The FMECA analysis is a “bottom up” analysis spanning the function failure modes of each subsystem. Some classical propulsion subsystem SPFs are shown in the FMECA snapshot in figure 2, these are common to N2H4 blow-down propulsion systems and are not unique to the SMAP mission.

Item				Failure			Effects Pre-Mitigation (e.g. no redundancy)		Mitigation		Criticality Post Mitigation	
Subsystem	Tier 1	Tier 2	Tier 3	Mode	Most Probable Cause	Phase of Criticality or Operational Mode of Equipment	Local/Subsystem	System	Failure Mode Detection	Compensating Provisions	Likelihood [1-3]	Impact Mission [1-8]
Propulsion												
	Propellant Storage											
		Propellant Tank										
			Shell / Tubing	External Leakage	Launch loads, fatigue	All	Uncontrolled loss of N2H4.	LOV: Loss OTM and RCS functions.	Pressure transducer detects pressure drop	Tank design margin > 2:1 on pressure	1	6
			Shell / Tubing	External Leakage	Material defect	All	Uncontrolled loss of N2H4.	LOV: Loss OTM and RCS functions.	Pressure transducer detects pressure drop	Tank is a facture critical pressure vessel and has passed full NDI / Acceptance testing and processing	1	6
			Shell / Tubing	External Leakage	Micrometeriod Damage	All	Uncontrolled loss of N2H4.	LOV: Loss OTM and RCS functions.	Pressure transducer detects pressure drop	Micrometeriod Sheilding	1	6
			Diaphragm	Internal Leakage across diaphragm	Damaged During Loading/Testing	All	Inability to deliver required N2H4	LOV: Loss OTM and RCS functions.	Loss of thrust before expected propellant depletion	Diaphragm integrity tested at launch facility just prior to propellant loading. Loading procedures preclude reverse diaphragm pressurization or excessive pressure loading	1	6

Figure 2: Snapshot of SMAP Functional FMECA

Follwing the generation of the FMECA and FTA, the leaves and branches of the FTA were combined with the failure modes in the system FMECA to comprise a comprehensive failure space. The failure space represents the collection of ways the SMAP hardware, functional behaviors and mission phases can fail. Note that the failure space is not a protected fault set, some of these failures modes are not mitigated and represent single point failures (SPF) of the SMAP mission. The failure space follows the same template as the Functional FMECA and is imported into a database program called FileMaker Pro. Duplicate failure modes (failures which are common to both the FMECA and FTA) are removed manually by fault protection engineers. Using this database, the fault protection team is then able to create relationships between the failure modes that have been identified and on-board compensating provisions such as FSW-based error monitors / responses or preventative measures such as a ground-based phasing test, and even specific verification and validation activities. The database can, and has, been used to inform risk trades given design changes.

III.A SMAP Risk Mitigation Matrix

The core of the SMAP fault management design is performed using a FileMaker Pro database in what we have called the SMAP Risk Mitigation Matrix. The mitigation matrix is an always-accessible online database that engineers can access and evaluate the current fault management design (as opposed to instanced documents or spreadsheets). The mitigation matrix refines upon the failure space to include preventative measures (tests, design features, analyses executed proactively during design do to reduce the likelihood of a particular failure mode occurring in flight) and mitigations (things which can be done reactively to reduce the consequence given a failure has occurred). Figure 3 provides a snapshot of the mitigation matrix; this particular failure mode is an inability to properly control the RWA speed due to a phasing error. The SMAP fault management design has both compensating provisions and preventative measures in place, which make this particular failure mode an acceptable risk. Ground based phasing tests performed on the flight vehicle reduce the likelihood of an RWA phasing failure (but not eliminate the possibility altogether), and on-board error monitors and responses detect the symptoms that result from such a failure (E.g. wheel speed control errors, spacecraft attitude and rate control errors, etc.) and place the observatory into a thruster based control mode which does not use the RWAs. Ground based recovery is required for resumption of science activities, which likely involves an update to the wheel transformation matrix.

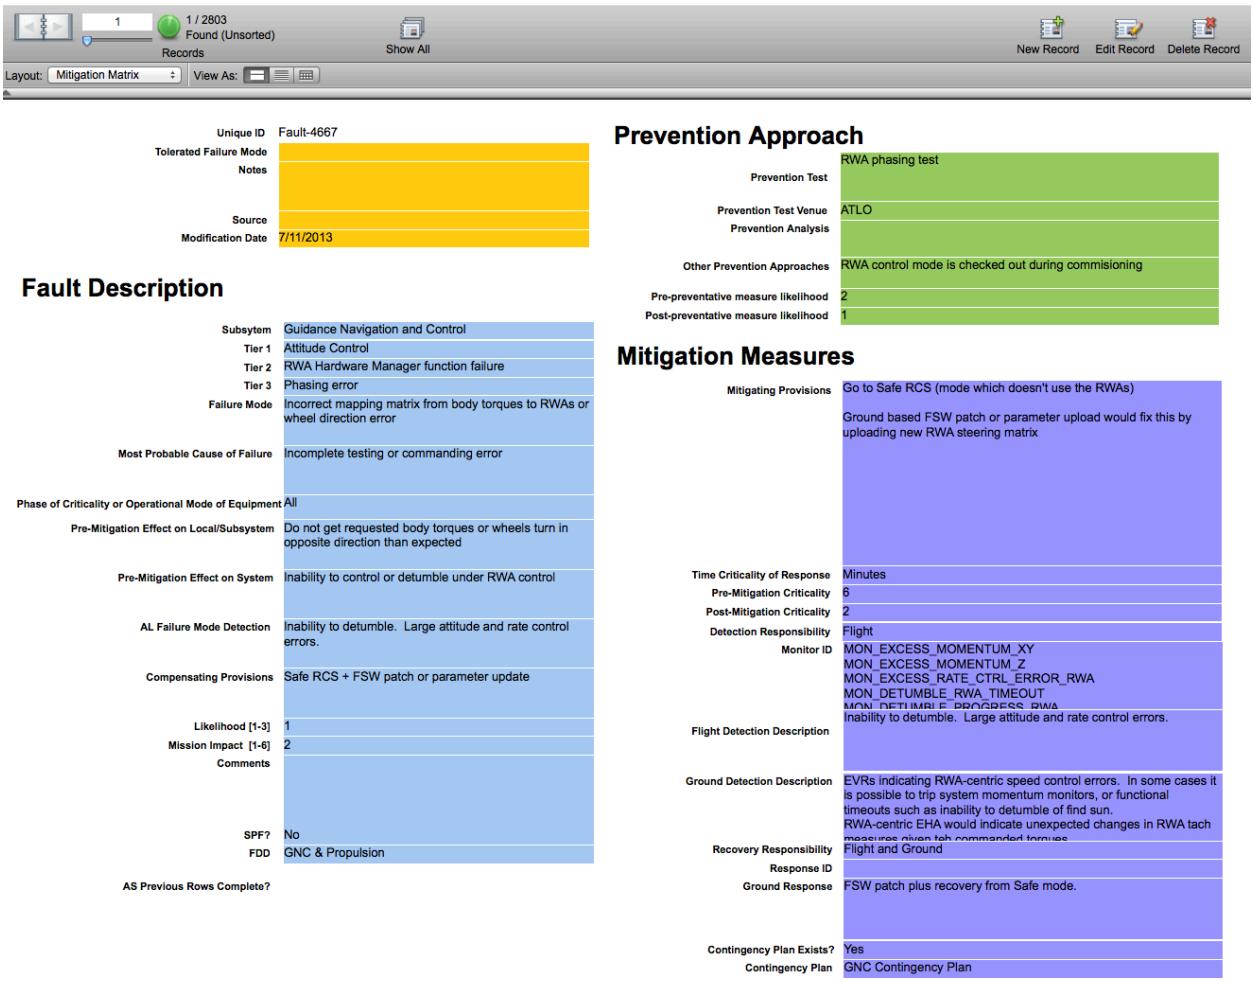

1 / 2803 Found (Unsorted) Records

Show All

New Record | Edit Record | Delete Record

Layout: Mitigation Matrix | View As:

Unique ID	Fault-4667
Tolerated Failure Mode	
Notes	
Source	
Modification Date	7/11/2013

Fault Description

Subsytem	Guidance Navigation and Control
Tier 1	Attitude Control
Tier 2	RWA Hardware Manager function failure
Tier 3	Phasing error
Failure Mode	Incorrect mapping matrix from body torques to RWAs or wheel direction error
Most Probable Cause of Failure	Incomplete testing or commanding error
Phase of Criticality or Operational Mode of Equipment	All
Pre-Mitigation Effect on Local/Subsystem	Do not get requested body torques or wheels turn in opposite direction than expected
Pre-Mitigation Effect on System	Inability to control or detumble under RWA control
AL Failure Mode Detection	Inability to detumble. Large attitude and rate control errors.
Compensating Provisions	Safe RCS + FSW patch or parameter update
Likelihood [1-3]	1
Mission Impact [1-6]	2
Comments	
SPF?	No
FDD	GNC & Propulsion
AS Previous Rows Complete?	

Prevention Approach

Prevention Test	RWA phasing test
Prevention Test Venue	ATLO
Prevention Analysis	
Other Prevention Approaches	RWA control mode is checked out during commisioning
Pre-preventative measure likelihood	2
Post-preventative measure likelihood	1

Mitigation Measures

Mitigating Provisions	Go to Safe RCS (mode which doesn't use the RWAs) Ground based FSW patch or parameter upload would fix this by uploading new RWA steering matrix
Time Criticality of Response	Minutes
Pre-Mitigation Criticality	6
Post-Mitigation Criticality	2
Detection Responsibility	Flight
Monitor ID	MON_EXCESS_MOMENTUM_XY MON_EXCESS_MOMENTUM_Z MON_EXCESS_RATE_CTRL_ERROR_RWA MON_DETUMBLE_RWA_TIMEOUT MON_DETUMBLE_PROGRESS_RWA
Flight Detection Description	Inability to detumble. Large attitude and rate control errors.
Ground Detection Description	EVRs indicating RWA-centric speed control errors. In some cases it is possible to trip system momentum monitors, or functional timeouts such as inability to detumble of find sun. RWA-centric EHA would indicate unexpected changes in RWA tach measures given teh commanded torques
Recovery Responsibility	Flight and Ground
Response ID	
Ground Response	FSW patch plus recovery from Safe mode.
Contingency Plan Exists?	Yes
Contingency Plan	GNC Contingency Plan

Figure 3: Snapshot of SMAP Risk Mitigation Matrix

The mitigation matrix contains a record for each failure mode in the failure space; only one record (out of thousands) is displayed in figure 3. Somewhat unique to the SMAP fault management process is that engineers evaluated for each failure mode the time to criticality (TTC) of response. This represents the time it takes for a failure effect (e.g. loss of attitude control due to a wheel phasing error) to take a mission critical effect (such as loss of vehicle due to battery depletion). The monitor ID field creates a relationship between the analysis and the existing error monitors (also contained in a database). Figure 4 shows the MM with portals to the monitor and response database based on the monitors listed. A portal is just a small view into another database, in this case the monitor and response database which shares linkages with the MM database based on the monitor ID field. The monitor and response database has a detailed layout providing much more information than is shown in Figure 4.

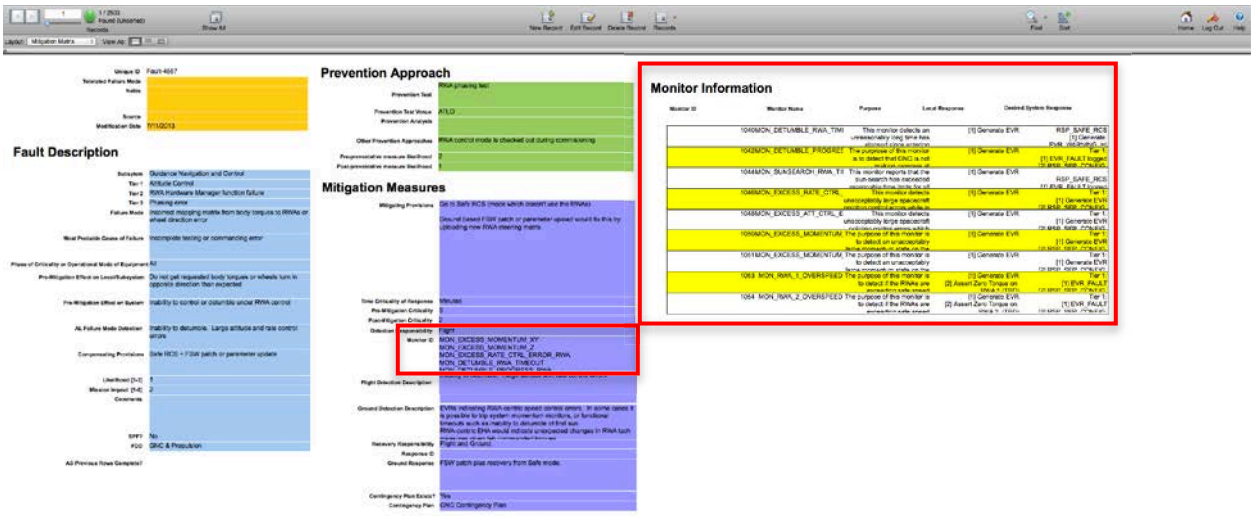

Figure 4: Mitigation Matrix with Portals to Error Monitor Database

The mitigation matrix is used to allocate fault management responsibilities to flight (on-board hardware and FSW-based algorithms) or ground operators. The mitigation matrix allocates the detection responsibilities (e.g. on-board FSW-based error monitors, ground operators, hardware algorithms, etc.) and corrective action / recovery responsibilities to the appropriate teams (E.g. FSW, hardware FP, ground operators). The recovery responsibilities are typically shared responsibilities between FSW algorithms and ground operators to resume science. The TTC for a failure mode is used as a first rough cut to help allocate flight and ground fault management responsibilities before negation between the flight subsystems (including FSW) and ground operators. For example failures which result in the loss of vehicle after tens of minutes (or hours) cannot be reasonably left to ground operators to detect since the observatory is not in continual contact with ground operators. When applicable, contingency plans are referenced in the mitigation matrix as well. As a general rule of thumb on SMAP, no fault management responsibilities with timeliness requirements less than one day are allocated to the ground operators, there are some limited exceptions during the launch phase.

In addition to the information contained within the mitigation matrix, the FP team conducted three additional analyses: timing analysis, interference analysis, and coverage analysis. These analyses were completed to facilitate the design specification prior to FSW implementation of the on-board FP design and similarly influence the ground-based fault management responsibilities. They answer systematic questions that confirm the intended coverage and effectiveness of the fault management design. Does the flight and ground FP design cover the expected/required failure space? Does the system respond to failures in a timely manner? Can the onboard FP design interfere with itself or other behaviors resulting in unacceptable or unintended interactions?

III.B Timing Analysis

The Timing analysis confirms that the FP design (both flight and ground) can detect and respond to observed errors in a timely manner. The timing analysis is meant to be used as a tool to validate that the FM design meets the timeliness requirements. While this is a relatively simplistic check early in the design process prior to FSW implementation, the intent is to help reduce the likelihood of finding problems later which could potentially be more costly in the future. The timing analysis evaluates whether or not the FP design can detect, contain, and mitigate the effects of a fault prior to the time to criticality being reached. The time to criticality is defined as the time it takes for the symptoms of a fault to reach a critical failure effect – resulting in the permanent inability of the system to acceptably meet requirements.

The first step of the timing analysis identifies the limited resource which is threatened by the failure mode, such as: science availability, ground communication, spacecraft power, thermal or physical damage to a device (such as an RWA over-speed condition damaging the RWA bearings). Based on the limiting resource, the time to criticality can be assessed. For example the loss of attitude control results in the inability to accomplish spacecraft pointing, but the time-limiting resource in this case is thermal and power considerations. Unsafe attitudes can result in overheating of some components and a power negative state. Note that "damage to device" is often essentially instantaneous, and reactive detection and response post failure is sometimes not able to meet timeliness requirements. In these cases the FP design proactively detects a condition *before* an unacceptable condition is reach

to allow sufficient time to detect and mitigate the failure. For example, the threshold for the RWA overspeed monitoring does not monitor at the limit at which damage is done to the wheel unit (WU), but rather at a limit below the level where damage is incurred. Analysis is then applied to ensure that the WU cannot increase to the level where damage is incurred from the point where the on-board FP detects the error condition. Next the FP engineers assess the time required for the autonomous (or ground) actions to complete. Note that the recovery time is not necessary the time when the last command is dispatched. In the event of an attitude control failure, it may take tens of minutes to bring the spacecraft attitude back to a safe state. Simulations of the attitude control performance and recovery timeline requirements are concatenated together to calculate the estimated recovery time. The timing analysis is considered passed if time to correct the failure is less than the time to criticality of the fault. A fault is not mitigated if the monitor and associated response takes long enough such that damage to hardware has already taken place. It is worth noting that the pass/fail criteria for timing analysis is not a pure recovery time < time to criticality equation. Engineering judgement is applied for each failure mode to evaluate if the design meets the timeliness requirements with sufficient margin. In some cases the fault scenario will be further evaluated during the formal V&V campaign.

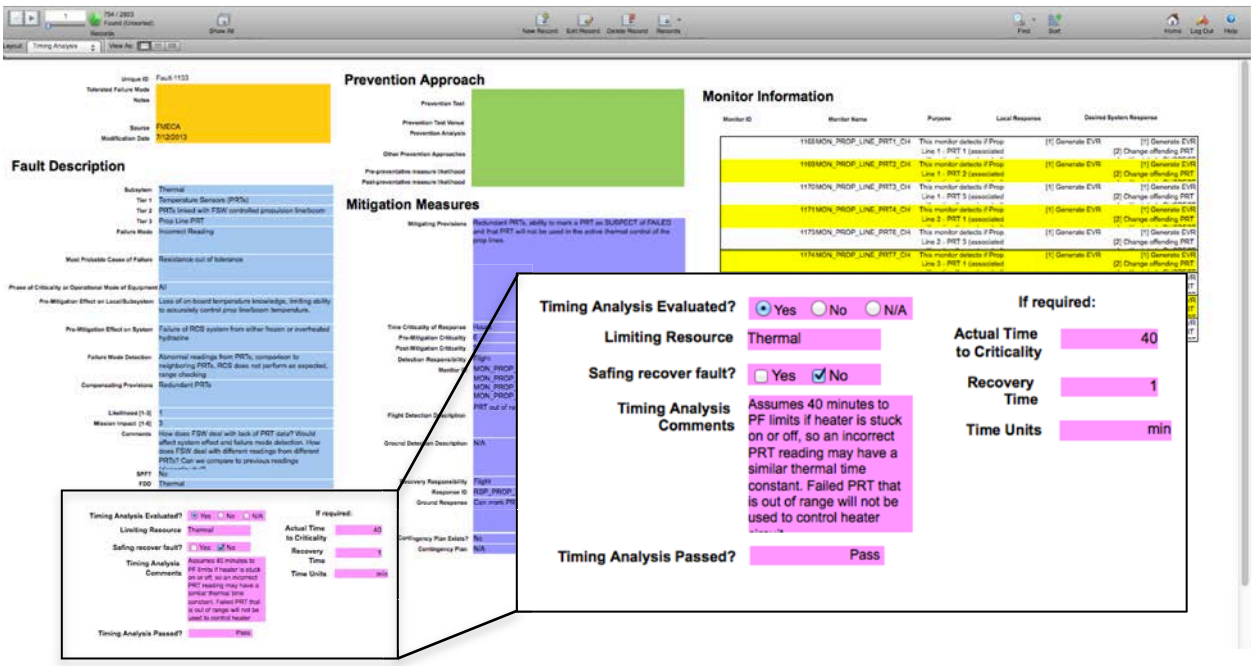

Figure 5 – Timing Analysis Example

The example in Figure 5 above shows the timing analysis for a failure of one of the three platinum resistence thermometers (PRTs) used by FSW to control an active thermal control zone. 3 PRTs are used in each thermal control zone, and SMAP is tolerant to a failure in any one measurement. Comparison checking is performed between the three sensor and a bad sensor measurement is excluded from the temperature solution. In the event of persistent bad measurements from a single sensor, the on-board FP will mark the the health of that unit "suspect" and it will not be used by the FSW.

III. C Interference Analysis

The goal of interference analysis is to evaluate if combinations of responses or monitors can interact in a manner that results in undesirable/indeterminate corrective actions or at worst preclude FP from achieving the proper end state required to mitigate a failure. The purpose is to evaluate the design and analyze potential conflicts which may arise before flight code is generated as a validation of the on-board error monitor / response design. This is not intended to be a substitute for testing, as many of these interactions are too complex or timing and state dependent to evaluate without simulation. Testing all combinatorics or states and timing conditions is an untenable scope and this analysis is used to identify potential areas of interference where targeting testing is required, and informs the fault management requirements and implementation.

The interference analysis was performed in two passes: Local Monitor/Response interaction and System Monitor/Response interaction (all tiers of every Response are evaluated). Some interference is precluded by functional characteristics of the Fault Protection design such as serial response execution, re-detection of errors following response completion or mode-dependent mapping of responses to declared errors. All system monitors and responses are placed into a large spreadsheet and an "N by N" interference analysis matrix is generated. Analysis is then performed on each of 36,864 combinations of system monitors and responses. There are three classification of interference between elements of FP:

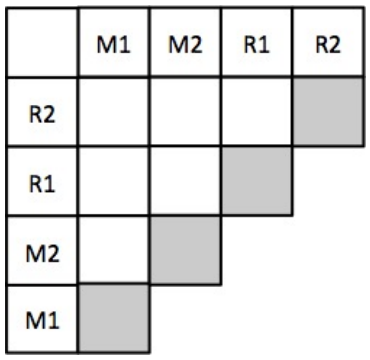

Figure 6 – Interference Matrix Layout

Response to response interference

Although responses are executed in series, there is still a potential for interference between responses. In response-to-response interference, the three key interference concerns are: one response "undo-ing" the actions of another response (e.g. over-temperature condition powers off SRU, can another response turn on that SRU and potentially cause a subsequent thermal failure), a response forces a transition which requires "higher" functionality (e.g. a scenario in which the FP design enters Safing and then later another response tries to place us back into Science), or after a combination of responses the system is left in an unsafe state (e.g. critical device is left unpowered – such as no powered IRU).

Monitor to response interference

In response-to-monitor interference, the interference being evaluated is if a monitor can/will be tripped due to a system response running. Several design changes were made to either the FP design or the low level FSW drivers to ensure monitors would not trip while devices were being reset. E.g. it is unacceptable if a GNC device reply timeout monitor trips every time the device is reset.

Monitor to monitor interference

In monitor-to-monitor interference, the interference being evaluated is if the local response of a monitor can trip another monitor, and also to ensure that if multiple monitors can be tripped as a result of a single fault that the end state is safe. When multiple monitors are tripped the possibility of multiple different responses being called is a concern which needs to be evaluated and is tested in the FP V&V campaign.

Systematically each pair is evaluated and dispositioned with one of following classifications.

Acceptable types of interference:

N: No interference due to analysis indicating functional separation or acceptable interference
M: No interference due to system mode-dependent features of the FM design
T: No interference due to hierarchy of fault detection or time separation of error detection
S: No interference due to serial response execution
A: No interference due to one of the activation rules
R: No interference due to the same response being invoked
O: No interference due to response timeout functionality
I: Interference exists by design (e.g. by design multiple monitors may trip)

Unacceptable types of interference

C: Potential conflict exists – more analysis / testing is required
D: Unclear whether or not undesirable interference will result, requires demonstration via testing

There were many instances of unacceptable interference were identified by the analysis and subsequently corrected via modifications to the FSW architecture or FP design. Types of modifications include: provisions within monitors and FSW to tolerate conditions resulting in transient loss of data, usage of response timeouts, tuning of thresholds and persistence limits, and new monitors and responses. One interesting finding discovered that it is unsafe to execute some corrective actions repeatedly. For example rapid repeated reassertion of some pulsed power switches can result in an unsafe thermal condition. It was discovered that the state assertion could be done by multiple un-correlated responses without knowledge of previous corrective actions taken by other responses and an

unsafe thermal condition can arise. Changes were made to ensure that only a single source can take these actions, and can throttle excessive attempts at recommending state.

III. D Coverage Analysis

The coverage analysis was performed to make sure that each failure mode that needs Flight detection is linked to at least one monitor. Additionally to ensure that all monitors have a purpose, analysis was performed that each monitor is linked to at least one failure mode in the Mitigation Matrix. One step further, the analysis also determines if all of the mitigating provisions required on the observatory take place either through the actions of flight responses or ground actions. The fault management design for spacecraft, even single-string earth orbitors, can be quite an involved process. This type of high level analysis ensures that the system is complete and that there are no extrenuous monitors created during the design phase.

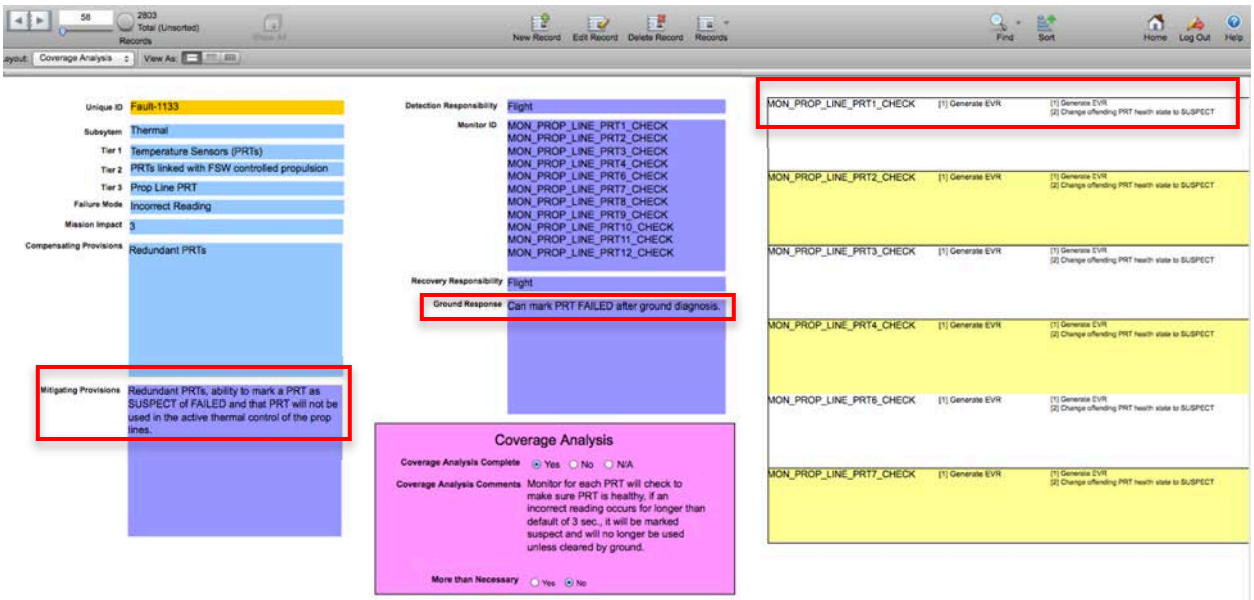

Figure 7 – Snapshot of Coverage Analysis Layout

To simplify the coverage analysis and to simplify the FP design, no provision is made if the coverage analysis indicates that the design takes more corrective actions than necessary. For example, in the event of an avionics failure which could be cleared by a reset of a single field-programmable gate array (FPGA), a reset of the entire avionics box is acceptable as a mitigation instead of just resetting the single FPGA. Engineering judgement and science availability calculations were performed on case-by-case basis to evaluate if the design was overly simplistic. Figure 7 shows the coverage analysis layout within the mitigation matrix.

IV. Conclusions and Future Work

The mitigation matrix and design analysis performed on SMAP resulted in many detailed design changes which are typically not considered at the early stages of the project life cycle. These analyses allowed the FP team to rigorously analyze the performance fo the system at an early stage in the project lifecycle. They provided an early design validation and resulted in a mature set of requirements being delivered to the FSW teams for implemention. Due to the timing that they were performed they have allowed for relatively fluid design changes without significant FSW rework or retesting. The coverage analysis was a particularily easy and important analysis to rigorously ensure that the FSW based FP really protects against the required failure space without having to verify each of thousands of low level fault scenarios during the testing campaign. The interference analysis and timing analysis have influenced much of the FSW implementation by early detection of potential areas of concern.

The mitigation matrix has played an important role in defining and coordinating the ground system fault management responsibilities. Ground FM responsibilities have been distilled from the mitigation matrix analysis in the form of requirements, contingency plans, flight rules, and trending/monitoring responsibilities. The fault

protection and flight operations teams have coordinated to document the expectations into a single source and listed the mechanisms that the ground operators can detect that a failure has occurred and also listed the appropriate steps that the operators need to take in response to the failure.

The initial purpose of the mitigation matrix is to inform the FM deisgn and requirements, the mitigation matrix has also allowed the team to make informed risk trades when considering addition or subtraction of particular deisgn elements. For example the mitigation matrix provides assessment of the criticality (impact) of failures with and without the on-board fault protection design. Figure 8 to the right indicates the criticality given failures with and without the on-board mitigations that are in place. In figure 8, a criticality of a "5" and "6" indicate loss of mission and loss of vehicle respectively. These failures are considered SFPs. A criticality of "4" indicates significant science degradation but that the minimum science floor can be accomplished. Criticality ratings of "3" indicate a loss of redundancy has occurred but all requirements can be satisfied. Ratings of "2" and "1" represent minor system impacts, a "2" indicates a loss of protective feature or a subsequent unrelated failure can result in loss of vehicle. Ignoring likelihood, the on-board FP design significantly reduces the system impact of many failure modes, this is true dispite the limited hardware redundancy on SMAP.

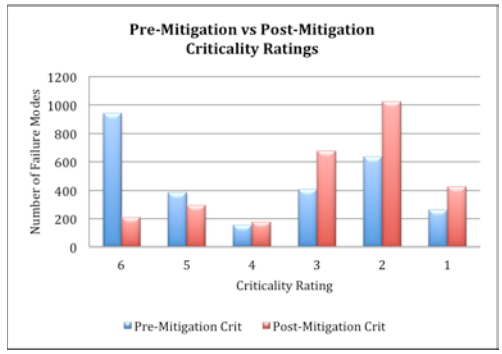

Figure 8 – Fault Criticality Pre-Mitigation and Post Mitigation

The SMAP FP design analyses described in this paper are not intended to replace formal testing of the system but rather to identify issues early in the detailed design process. In addition to supporting the detailed design requirements generation, the analyses have also provided insight where testing may be the most beneficial. For example, analysis has demonstrated that the accumulation of system momentum as a result of control failures during propulsive maneuvers is a stressful scenario from a timing perspective. Momentum can quickly accumulate in the event of some propulsion system failures and subsequent slewing under RWA control would requires large gyroscopic torques. In order to allow for RWA-based recovery, the FP design must quickly detect propulsion failures and stop the propulsive maneuver before an excessive momentum condition can arise. The team has identified several key scenarios with stressful timing requirements or difficult to analyze interference with the rest of the system and placed these into the planned FP-centric testing which is controlled in a database linked to the mitigation matrix.

Although the work performed by the SMAP FP team has improved the design requirement process by understanding and identifying areas of concern at an early stage, there is much room for improvement. The analyses performed by the SMAP FP team represent an instantaneous design point based on the FMECAs and FTAs, and require manual updates in response to spacecraft design changes. Changes can and have been made with spot checking updates to the analyses. Careful up-front planning when performing FMECAs and FTAs in early mission design phases will help to automate the analyses performed by team members during detailed design generation and implementation. Of particular note, the interference analysis was performed manually and is very difficult to reproduce given a significant number of changes. The team recommends that the interference analysis be performed last when the design is relatively mature from other sources of design changes.

While very beneficial, the design analyses are rather simplistic in their assumptions. While this was a necessity given finite resources, it is possible to get a false positive sense of the fault tolerance without off-nominal scenario testing. For example the timing analysis assumes a single TTC and a single "tall tent-pole" limiting resource. In reality a spacecraft is a complex system of interactions and critical failure effects may not be realized in our simply modeling. At the instance that the analysis is performed the shortest TTC may indicate the restoration of active thermal control of a heater (controlled by FSW) as the limiting resource but the subtle change in an attitude control gain may result in the inability to achieve an acceptable power positive attitude and result in unacceptably low energy states of the battery.

Of particular interest for future work would be to rigorously model the FMECAs and FTAs, which could be used to show dependencies between faults and the subsequent detectable errors and symptoms. This could be relatively easily incorporated into the mitigation matrix work. Coupling the FMECA and FTA models with an executable model of the FSW behavior would allow for continual and potentially automated design analyses of significant benefit. A light weight model of the FSW interactions and timing coupled with these analyses would allow for real-time evaluation of the FM system and would constitute a significant improvement.

A key lesson learned by the SMAP FP effort is that some of the analyses performed occurred too late in the design process after the hardware requirements were set. Ideally the mitigation matrix analysis would occur prior to

hardware requirement generation and help to determine the adequacy of the fault tolerance intrinsic in the hardware. Some failures simply have a time to criticality that is fast enough to preclude reactive flight software or ground operator based mitigation. These types of failures require preventative measures in the hardware to preclude their very occurrence.

Overall the design analyses performed by the SMAP FP team have improved the FM design process. It is possible to make more informed trades regarding the fault tolerance of the observatory and rigorously check that the design assumptions which engineers make during early FMECA/FTA activities are implemented per their expectations. While much work can be done to improve these analyses for future missions, they represent a step in the right direction toward an analytical assessment of the FM deisgn.

Acknowledgments

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. The authors are grateful to many people, without whose help and inputs, this paper could not be written: John Day, Julie Wertz Chen, among others.

References

[1]Spencer, M., Wheeler, K., Chan, S., Piepmeier, J., Hudson, D., and Medeiros, J., "The Soil Moisture Active Passive (SMAP) Mission L-band Radar/Radiometer Instrument", Geocience and Remote Sensing Symposium (IGARSS), 2011 IEEE International, Honolulu, HI, July 25-30, 2010.

[2]Johnson S. and Day. J., "System Health Management Theory and Design Strategies", Infotech@Aerospace 2011, Louis, MI, March 29-31, 2011.