The Soil Moisture Active Passive Mission: Fault Protection Performance and Lessons Learned

Jessica Juneau Clark
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
626-720-1989
Jessica.Juneau@jpl.nasa.gov

Peter Meakin
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-393-5530
Peter.C.Meakin@jpl.nasa.gov

Abstract—**Fault protection as a discipline involves a collection of flight software logic and operational processes for detecting unacceptable anomalous behavior, responding prior to reaching criticality, restricting the propagation of a failure beyond a fault containment region, and recovering the vehicle back to full or degraded functionality if possible. The System Fault Protection (SFP) design for the SMAP Earth orbiter was put to the test during its 90-day vehicle Commissioning activities. During this time, the SFP software autonomously protected the vehicle from multiple faults to critical hardware, and the operations team successfully returned the observatory to its science state. The SFP also performed well in the presence of anomalous behavior below true safety limits by not taking unnecessary response actions, instead allowing the operations team time to monitor the behavior. Certain aspects of the SFP design were modified during operations via both parameter updates and a full flight software update in order to better match the vehicle behavior in the flight environment. An evaluation of the SMAP SFP performance during vehicle Commissioning will be provided in this paper, as well as a set of lessons learned largely focused on visibility, SFP mutability in operations, responses to peripheral device faults, and Safe Mode recovery and design. By capturing some of the knowledge gained during SMAP Commissioning, it is intended that this paper provide guidance for making future System Fault Protection designs more robust and supportive of operations.**

TABLE OF CONTENTS

INTRODUCTION

On January 31, 2015, the Soil Moisture Active Passive (SMAP) satellite lifted off from Vandenberg Air Force Base inside of a Delta-II Heavy launch vehicle. The Earth orbiter entered a roughly 685-km, near-polar, sun-synchronous orbit in order to complete its 3-year baseline science mission: increasing the understanding of global processes that link the water, energy, and carbon cycles, in addition to increasing the capabilities of climate and weather prediction models. Using an on-board radiometer, synthetic aperture radar, and rotating 6-m deployable mesh antenna, SMAP was designed to provide global measurements of soil moisture and its freeze/thaw states. The SMAP Flight System includes engineering subsystems necessary to support the operation of the spacecraft and instrument, with the Mission System on the ground providing the personnel, processes, and equipment to support the prime mission. Figure 1 shows a conceptual view of SMAP in orbit with its deployed and spinning Reflector Boom Assembly (RBA).

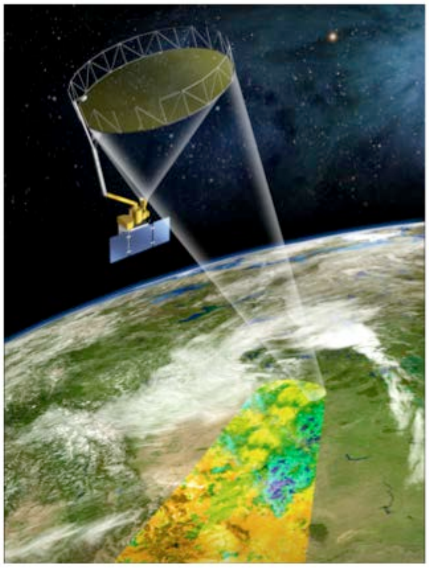

Figure 1. SMAP observatory artist's concept

In July of 2015, the SMAP radar experienced an anomaly with its High-Powered Amplifier (HPA) after which it stopped transmitting. After substantial investigation and recovery attempts from the team, it was determined by the mission that the radar could no longer produce data. While this results in some impact to SMAP science, the observatory continues to operate its radiometer instrument to produce high-quality soil moisture and freeze/thaw data. This paper will not focus on the radar HPA anomaly, as this did not occur during vehicle Commissioning, nor were there autonomous SMAP system Fault Protection responses for radar hardware failures.

To meet the mission fault tolerance requirements, a set of on-board system fault protection software exists to provide autonomous response to certain vehicle anomalies. The purpose of this paper is to provide an assessment of the SMAP fault protection performance during vehicle Commissioning as well as key design lessons learned. An overview of the SMAP FP architecture will also be provided for context and background. This paper does not intend to suggest a complete SFP architecture, as this is largely dependent upon mission-specific factors. Many of the SFP-specific lessons learned are relevant to similar, Earth orbiter missions with non-unique science data. Other insights have a wider range of applicability to missions with strict fault tolerance requirements including the importance of testbed fidelity, fault protection strategies for faults in non-critical hardware, and sufficient visibility into anomalous behaviors. By capturing some of the knowledge gained during SMAP Commissioning, it is intended that this paper provide guidance for making future System Fault Protection designs more robust and supportive of operations.

1. FAULT PROTECTION DESIGN OVERVIEW

System Fault Protection (SFP or FP) as a discipline involves both the ground and flight systems in a collection of logic and processes for detecting unacceptable anomalous behavior, responding prior to reaching criticality, restricting the propagation of a failure beyond a fault containment region, and recovering the vehicle back to full or degraded functionality if possible. The SMAP on-board System Fault Protection uses a series of monitors to detect error conditions and to report them to a fault protection engine. Within the FP engine, the monitors are mapped to a given response that is then placed in a queue and evaluated by the engine based upon a set of activation rules. Queued responses are either thrown out of the engine or executed in a particular order depending upon the built-in activation rules. Engine queuing for SMAP occurs in a serial manner so as to prevent unintended response interactions, such as un-doing a device swap. Figure 2 provides a high-level diagram of the SMAP System Fault Protection architecture with regards to response, monitor, and FP engine interactions.

Fault protection error monitors are designed to detect deviations from nominal performance by performing a repeated error test. For a majority of SMAP FP monitors, error tests are evaluated with respect to a given threshold parameter (maximum value or minimum acceptable value) and a persistence (maximum allowable duration prior to monitor reporting an error). Some monitors, such as timeouts, have only a persistence parameter. When the parameterized threshold value is violated for the duration of the persistence, the SFP monitor declares an error to the SFP engine, or "trips." A cartoon demonstrating flight telemetry that has violated it persistence and threshold limits and tripped FP is shown in Figure 3.

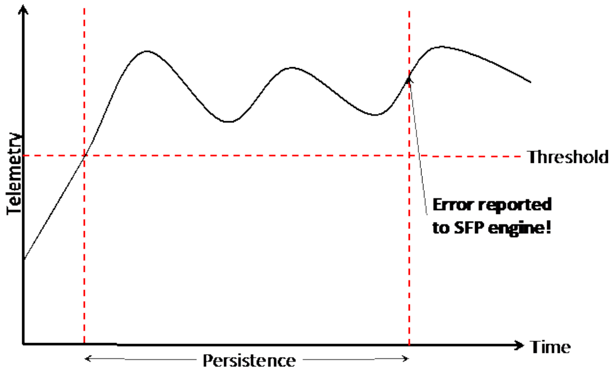

Figure 3. Example of behavior "tripping" FP

In certain scenarios, fault protection monitors can be disabled, or masked, to prevent unintended interactions with other SFP responses, to prevent false-trips, or to remove obsolete monitors. In addition to being "Masked" or "Unmasked," SMAP fault protection monitors also have a monitor state that is reported in vehicle telemetry. A monitor whose error condition is currently being evaluated and whose threshold and persistence limits have not been exceeded is reported as "green." A monitor for whom both the threshold and persistence limits have been exceeded and who has not been "cleared" by the SFP engine is reported as "red." Finally, a monitor whose error condition cannot be evaluated (e.g. the device is powered off) is reported as "black" and cannot initiate SFP actions.

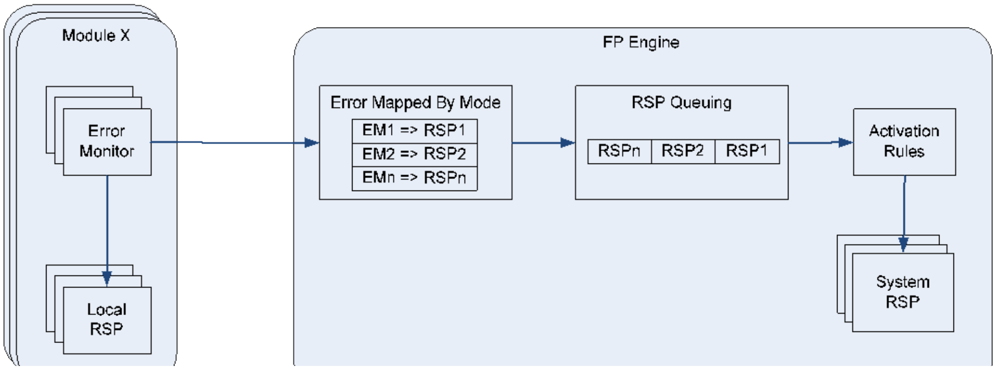

Figure 2. High-level diagram of SMAP SFP architecture.

Fault protection responses contain a set of actions performed when that particular response is "activated" by SFP. SMAP chose to use a tiered response architecture, where for multiple occurrences of a given monitor error declaration, different sets of actions are taken. These tiers typically increase in terms of system impact and are meant to start with the closest fault containment region and propagate outward if unsuccessful. For example, if Guidance Navigation and Control (GNC) attitude control errors were unacceptable, the first tier of SFP might power cycle the IMU (Inertial Measurement Unit). If that failed to correct the error, the second tier of FP could swap to the redundant IMU; and a final tier could power cycle the flight computer to clear out some avionics faults. In addition to masking SMAP FP monitors, the team can also choose to mask individual responses. This can be helpful in situations where the ground wants to avoid taking certain actions, like swapping to a redundant unit, regardless of the monitor that calls it.

For governing the overall system behavior, SMAP chose to use a state machine architecture for each of the major system modes. Upon entrance to each of the modes, a configurable table enforces certain device power states. This behavior is heavily integrated with SFP behavior as it allows for a known, enforced state upon entry into the vehicle safe modes. The current system mode is also evaluated in the SFP engine when activating SFP responses. The team can select to "map" SFP responses only to certain system modes, thus preventing monitors from triggering SFP actions in modes where it is undesirable.

As will be discussed later on in the Operations Lessons Learned section, SMAP chose to use system modes as different levels of safing. A "Standby" mode, in which the vehicle remains spinning and in the science attitude, was used for a subset of faults on hardware that is non-essential for maintaining the science configuration. For more substantial faults, SFP has a "Safe Mode" that leaves the vehicle spinning and using the reaction wheels, but places it in a sun-pointed rotisserie mode. Finally, for faults involving reaction wheels and spin hardware, or for lower level response tiers, SFP uses another "Safe Mode" that turns off the spin electronics to spin down the vehicle and switches to thruster control. As a general rule, the SFP design attempts to protect hardware that is non-critical to Safe Mode, typically by powering it off, whereas it attempts to recover hardware that is critical for Safe Mode (by power cycling, swapping to redundant units, etc…).

Entering one of the vehicle Safe Modes powers off both instruments, and in some cases spins down the vehicle. During the SFP design phase, entering Safe Mode had to be traded against impacts to the mission science timeline as a result of the Safe Mode implications. It was decided to weight SFP responses on the less forgiving side and safe the vehicle for most anomalous conditions. This decision was largely made due to a low expected frequency of vehicle anomalies. The FP team also wanted to more heavily rely upon the vehicle Safe Modes to enforce a well-tested state that relied upon lower-level system functionality, particularly during first time activities where system behavior in the presence of anomalies was not fully understood.

As a note, this paper will not focus on avionics reboot logic and reset dead-ending and will not describe individual fault protection monitor and response behaviors in detail.

OPERATIONAL PERFORMANCE

The performance of the SMAP System Fault Protection has been excellent in preserving critical functionality throughout observatory Commissioning in the presence of several anomalies. The on-board Flight Software (FSW)-based fault protection responded appropriately to several hardware-centric anomalies, protecting both the device health and enforcing a safe state that maintains health-critical functionality. These situations would have resulted in loss of the vehicle without an autonomous, on-board response. Additionally, the on-board SFP appropriately tolerated abnormal conditions that did not pose an immediate threat to vehicle health.

The first FP-initiated safing event was due to an unexpected reset of GNC hardware needed for stellar attitude estimation. Without stellar attitude updates, the attitude solution can be propagated based on rate measurements but not for indefinite periods of time. An on-board device communication FP monitor detected the anomalous device behavior and invoked a safing response, powering off the misbehaving device as opposed to commanding the device back to an operations state. The chosen response behavior is typical for devices on 1553 buses, as logical faults can propagate to other remote terminals on the bus and since most loss of bus communication failures result in no further 1553 data until autonomous action is taken. Although this protocol FP took a system level response that is typically necessary to protect device health and establish a safe state, it lost important diagnostic information by powering off the device quickly despite only a temporary communication outage during an unexpected reset. In response to this anomaly the operators recovered fully and also loosened the FP limits for the protocol monitor based on suspected root-cause of the hardware anomaly and analysis to confirm no vulnerabilities associated with the change.

The operations team's foresight to loosen the limits for the device protocol FP proved to be a wise decision as the hardware anomaly repeated itself again. This time the anomaly occurred over a communications pass and with the newly lengthened response time and closed-loop ground commanding actions, the operators were able to recover the device without entering Safe Mode or affecting the overall system state. A FSW update was eventually developed to modify the FP response to first recover the functionality of the GNC hardware without the need for ground intervention. In the new design, persistent anomalies would eventually result in the original safing response. Although this was a break from the simple architecture described in Section 2, the multiple occurrences of this fault convinced the operators that the anomaly was not as low probability as initially believed. This FSW update was successfully applied in flight, and shortly after the update the hardware anomaly was observed again. With the updated FP logic the on-board FSW collected necessary diagnostic data for ground operators, repaired the mis-behaving device, and ultimately returned it to a fully-functional state. The hardware anomaly has now occurred four times since launch. Due to the configurability of the FP control parameters and the flexibility to update FSW responses, the operations team has precluded this unexpectedly frequent anomaly from safing the vehicle and subsequently impacting the mission science timeline.

The second FP-initiated safing event was also due to an unexpected fault in GNC hardware related to the rate estimation of the vehicle. On-board FP detected the mis-behaving device and power cycled it, returning it to a fully functional state. In this case the on-board response also invoked a safing response as was necessary by FSW design to re-establish closed-loop attitude control. A second update to the FP design has since been designed to ensure a more elegant response to this fault, and the operations team is ready to proceed in the event that this anomaly proves to be a frequent occurrence.

In addition to responding appropriately to health-critical faults, the on-board FP also demonstrated the ability to tolerate conditions that did not directly threaten observatory health by not safing the vehicle. For example, several devices at times reported temperatures that were considered anomalous by the subsystems yet below hardware safety limits, and so FP did not respond, allowing the ground time to monitor and assess the situation. FP limits were also sufficiently high so as not to falsely trip during transient and expected “abnormal” conditions that occurred during device power on and off. For example, although radio power on events prior to communications passes resulted in brief configuration errors due to software timing, fault protection limits were set sufficiently high so as to “ride through” the expected transient behavior.

The key takeaway in all of these areas is that off-nominal events will occur in flight. The system design, including the SFP approach needs to protect against unexpected environmental effects as well as unexpected hardware and software performance. Functional “safety net” FP monitors play an important role in ensuring the health preservation of the SMAP spacecraft, whereas some of the device protocol FP limits were set too tight and required adjustment. Additionally it is important to ensure that FP protects against threats to the health of the observatory as opposed to responding to every unexpected or anomalous condition.

OPERATIONS LESSONS LEARNED

This section details key fault protection lessons learned during SMAP Commissioning in operations. Information is largely focused on FP design lessons learned during operations, and is grouped into the following broader categories:

- Visibility
- Fault Protection Mutability in Operations
- Robustness to Peripheral Device Faults, and
- Safe Mode Design and Recovery.

Visibility

Having sufficient visibility into anomalous behavior, System Fault Protection actions, and long-term FP statistics is essential for spacecraft operations. At a minimum, the operations team should be capable of quickly answering the following questions:

- Has System Fault Protection executed, and is the vehicle currently in Safe Mode?
- If so, what SFP monitors tripped and which response actions were taken?
- Have any devices been marked unhealthy by SFP?
- Has SFP swapped to a redundant unit as part of its response actions?
- Has the flight computer undergone a reset, clearing some of the SFP response history in volatile memory?
- If the vehicle state is nominal, are there any SFP monitors that are close (>80%) to tripping based upon trending data?

One of the most useful long-term and daily SFP trending efforts involves tracking actual monitor persistence counts and threshold spikes with respect to parameterized limits. This information is important for catching monitors that are close to tripping due to actual faults or due to “false positive” situations, which require persistence or threshold management in order to prevent unintentional SFP actions.

For SMAP, the number of counts on a monitor’s persistence is reported in spacecraft telemetry, but for ease of visualizing the maximum value reached over a period of time, the team typically used a high-water mark data product that was operationally telemetered every day. This was helpful both for trending and anomaly investigations. For example, by capturing and clearing high-water mark data daily, the fault protection team noticed that a protocol monitor for one of the radios was reaching up to 80% of its persistence limit. Upon investigation, it was discovered that these peaks occurred during radio power on events and were later confirmed to be within the realm of expected outages during power on per the vendor – the team had simply not observed this during testing with the flight unit. As a result, the monitor persistence parameter was increased to a new safety limit to prevent unintentional SFP action. Other times, SFP persistence counts notified the team of real device anomalies. When one of the GNC devices critical for maintaining SMAP's science attitude unexpectedly reset, an SFP protocol monitor for that device tripped and safed the vehicle. Later on, the decision was made to increase the persistence on that device's protocol monitor and add new SFP behavior to respond to unexpected device resets. Rationale for the software change is discussed in the subsection on Robustness to Peripheral Device Faults. For future resets, the high-water mark data provided the ops team with device outage durations for trending of reset behavior.

A lack of channelized telemetry containing monitor error condition values with respect to threshold limits created some visibility challenges for the SMAP FP team. Because these telemetry items were not explicitly specified for flight software developers, some threshold high water mark data during operations remained unknown. Some values required ground calculation, and fortunately others were already reported by corresponding flight software modules. Threshold high water mark information aids in catching near-misses and anomalous behavior below the chosen FP limit. For full visibility into SFP behavior, it is suggested that the team have telemetered error condition values versus monitor thresholds reported directly by the SFP monitors.

Another option for providing some insight into monitor error evaluation, which SMAP chose not to utilize, is to have an intermediate "yellow" monitor state that is reported in spacecraft telemetry. As opposed to reporting a monitor as "green" until the threshold has been exceeded for the persistence duration, the monitor is reported as "yellow" if just the threshold is exceeded. This will not catch near-misses, where a large percentage of the threshold (e.g. 90%) is reached, but it does raise a flag to the operations team that a critical threshold has been exceeded. Figure 4 shows an example of monitor "green," "red," and "yellow" states for a given spacecraft state history (e.g. a device input current measurement).

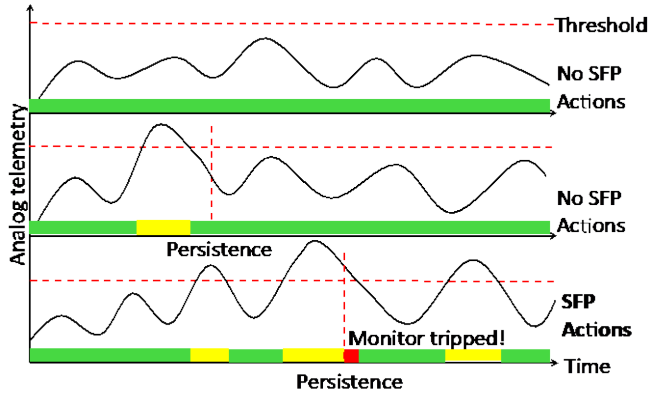

Figure 4. Monitor state, including "yellow" state for given telemetry

A complement to SFP spacecraft telemetry is the ground alarm. These are checks built into the ground software with limits managed by the ops team to alert them to anomalous behavior. These limits can be set to values lower than the SFP thresholds so as to provide an early warning that a telemetered item is in undesirable territory. In addition, the ground alarms for SMAP were built into a system that would automatically alert the team via text message if a "red" (critical) alarm tripped. With the automatic notification capabilities, the red alarms were used to also notify the team as soon as downlinked telemetry indicated SFP had executed safing.

A final, highly recommended tool for operations visibility into SFP state is a color-coded dashboard. The SMAP FP ops team utilized a dashboard that contained a grid of boxes, with each box mapping to an SFP telemetry channel of interest – for example, a monitor state or a device health state. If a box in the grid is red, it notifies the operations team that a channel is in an undesired state. This provides a quick-look assessment of hundreds of SFP telemetry channels all at once.

Fault Protection Mutability in Operations

The performance of the SMAP System Fault Protection is heavily tied to factors such as the system modes state machine and hardware configuration enforcement, on-board timing and telemetry sampling, and expected ground-side recovery actions. Because of this, the fault protection pre-launch Verification and Validation (V&V) can be a challenge that leaves some aspects of validation to observations in the actual flight environment. To account for these uncertainties, it is important to implement a System Fault Protection design that can be easily modified during operations. However, with a large number of SFP "knobs" comes the difficulty of tracking vehicle state and in matching vehicle state in ground testing. Therefore, a balance is needed between configurable and hard-coded aspects of the SFP design.

For SMAP, the decision was made to parameterize monitor threshold and persistence limits, allowing the ground to easily change them without a flight software update. The same is true for device health states, the SFP error count (number of monitors that have tripped), and the response tier counters, but those characteristics were expected to be changed as part of Safe Mode or SFP recovery following an anomaly and not to just change the behavior of SFP. The mutability of SFP thresholds and persistence was incredibly valuable, and those parameters were tuned many times during Commissioning as the team gained a better understanding of flight system behavior in the real environment.

On the other hand, many aspects of SFP were hard-coded in the flight software in order to simplify parameter management as well as preventing unsafe changes to fault protection. SFP monitor error checking logic as well as response actions were originally all hard-coded – and this was important to prevent haphazard operational changes to thoroughly tested software.

However, more flexibility in monitor logic and response behavior was added into a flight software update, which was largely implemented to adjust the SFP response to an unexpected device reset. In the new build of flight software, an SFP monitor and response were added that contained mutable global variables because of uncertainties in the device reset behavior. Having this flexibility was critical, because there was no EM for ground testing of the device's reset behavior.

Global variables or parameters could also be used for bulk disabling of FP monitors after a critical event completion. This was not used as part of the SMAP SFP design, but would have simplified some of the "cleanup" activities following one-time events, like reflector and boom deployment. Having a Boolean for "deployment complete," for example, would have allowed for masking of all now-obsolete fault protection monitors so that they do not falsely trip in the future. However, the addition of global variables must be traded against impacts to operations efficiency, given that for every new SFP variable or parameter comes the onus of managing them.

The state machine system modes architecture also provided additional flexibility for managing SFP behavior. Because the system modes uses tables to reinforce certain hardware states, system modes can be repurposed for post-Commissioning ops scenarios. For example, although SMAP system modes existed for Reflector Boom Assembly (RBA) deployment and initial spin up, these modes could also be modified for off-nominal scenarios. In the RBA deployment mode, the GNC goes into an "Idle" state which can also be utilized in situations that would otherwise unintentionally trip fault protection. The spin up mode can also be utilized for spinning to a new spin rate due to unexpected spacecraft wobble, all without requiring an update to flight software. In addition, because SMAP fault protection responses could be mapped only to certain system modes, obsolete monitors (like RBA deployment-specific monitors) could be disabled by modifying the mapping. Having the ability to repurpose the state machine modes is a valuable tool for both Commissioning "cleanup" as well as anomaly recovery.

Because changes to fault protection are likely to occur in operations, especially during vehicle commissioning, it is highly suggested that sufficient FP engineers be staffed in order to support parameter changes and flight software updates. During SMAP Commissioning, 3 FP engineers were staffed to support day-to-day operations, but many more were recalled back to the project when flight software changes became necessary. The availability of past FP team members was essential to completing a timely and thorough design, V&V, and implementation of the new flight software build. Documentation and maintenance of shared repositories from pre-launch verification testing was also essential in reducing time for test script development.

In summary:

- Having appropriate aspects of the SFP design that can be easily modified by the operations team is a powerful capability.

- This flexibility is especially valuable during spacecraft commissioning and first-time events.

- However, such flexibility should be considered alongside the challenges of maintaining ground knowledge of vehicle state and the testing heritage of certain software.

Robustness to Peripheral Device Faults

The SMAP System Fault Protection response architecture for a peripheral device fault became heavily vetted during Commissioning. Multiple times, supporting GNC hardware experienced unexpected resets, believed to be due to the radiation environment. For all occurrences device protocol SFP tripped and performed both local actions (e.g. power cycling the device) and system actions to place the vehicle into Safe Mode. As discussed in the Operational Performance section, a fault protection response was absolutely necessary in the presence of these anomalies, as the devices provided critical inputs to the GNC algorithms given the current GNC submode. SFP responses to device protocol faults were also intentionally designed to take a larger, system response rather than attempt a local recovery, especially for hardware whose anomalous behaviors were not as well understood via ground testing or previous flight data.

Though the original SFP design helped to recover the vehicle and placed it into a known safe state, the frequency of the device resets and the timeline hits to mission science prompted an update to the fault protection software. For this flight software update, the desire was to have local repair actions on the device for the first two unexpected resets and then to place the vehicle in Safe Mode on the third. Choosing the right settings for the local repair actions presented a challenge, as the unexpected device resets had only been observed in flight and the team had no Engineering Model (EM) with which to test the response on the ground. Earlier during the project development cycle, the decision was made to not purchase an EM of the device due to multiple factors including the associated cost and the expected similarity to previously flown models. However, during operations the device exhibited unexpected behavior that was not captured in the testbed software simulation and that could not be fully understood without the actual hardware and firmware present. For future projects, having access to Engineering Models of complex devices during operations will provide the test venue fidelity necessary to support anomaly investigations.

Without sufficient test venue fidelity, the team had to use the smaller set of resets observed in flight and build that observed behavior into the simulation software to support ground testing. Some of the reset behavior remained unknown, as fault protection power cycles or powers off misbehaving devices, resulting in lost volatile data. Therefore, the team added global variables in the new response in order to account for the reset signature uncertainties and to allow the team to "tune" the responses as necessary in ops.

Having a tiered response architecture also proved valuable in implementing the in-flight software changes. Because the team was uncertain as to how many local recovery attempts would be sufficient, multiple tiers were added to the new response in order to repeat the same actions prior to safing. The new response tiers provided a "second chance" for the device to recover with just a local response, while still guaranteeing that the vehicle would enter Safe mode upon the third reset. Figure 5 provides a high-level example of a both a flexible and more rigid monitor and response architecture for a peripheral device fault.

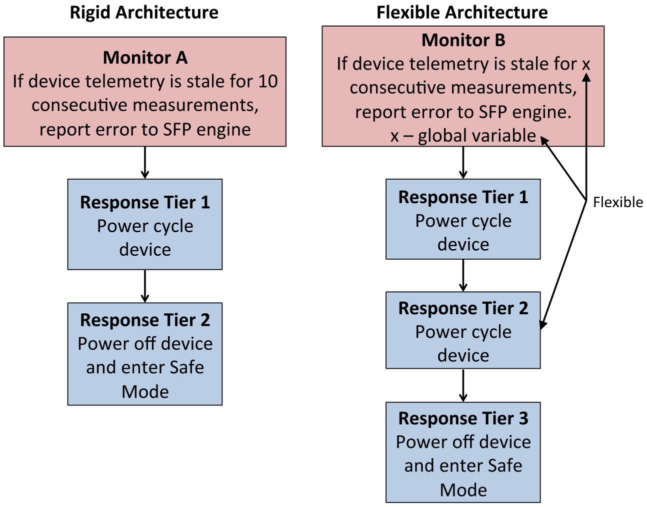

Rigid Architecture

Monitor A
If device telemetry is stale for 10 consecutive measurements, report error to SFP engine

Response Tier 1
Power cycle device

Response Tier 2
Power off device and enter Safe Mode

Flexible Architecture

Monitor B
If device telemetry is stale for x consecutive measurements, report error to SFP engine.
x – global variable

Response Tier 1
Power cycle device

Response Tier 2
Power cycle device

Response Tier 3
Power off device and enter Safe Mode

Figure 5. Rigid and flexible monitor/response architecture.

Overall, given the flexibility in the fault protection architecture (with a tiered response structure) and the allowance of ground tweaks of monitor error checks (via global variables), the team was able to implement an effective FP solution to the unexpected device resets despite lacking extensive ground test data to properly characterize the behavior.

Safe Mode Design and Recovery

The purpose of a vehicle "safe mode" is to provide a known state that is power-positive, comm-positive, and thermally safe – a mode in which the vehicle can operate for extended durations while the operations team completes its anomaly response. For SMAP, multiple modes in the system modes state machine were utilized as different types and "levels" of safing.

For a subset of faults, especially on hardware that is non-essential to maintaining spun-side control and the science attitude, a "Standby" mode was used as a less severe SFP end-state. As opposed to the vehicle Safe Modes, the "Standby" mode allowed the vehicle to remain in its science attitude with the spin-side still spinning. It utilized the system modes configuration tables to enforce certain hardware states and relied upon built-in response actions that were specific to the expected fault. During the SFP design phase, response actions were biased toward utilizing the "bigger hammer" Safe Modes given SMAP's non-unique science data requirements. However, at later points in SMAP Commissioning, the team strongly considered modifying some fault protection responses to transition to "Standby" instead of one of the vehicle Safe Modes as part of its first tier to increase project science efficiency. For faults where a delayed full-safing was acceptable, this would allow for a lower impact to science and anomaly recovery, as the instruments could remain powered on, GNC hardware would not require reconfiguration, and the vehicle spun-side would not have to be spun up back to the science rate.

For more severe faults, especially those involving critical GNC actuators and sensors or spin hardware, SFP utilizes the less-forgiving Safe Modes to guarantee a safe end state. One of the Safe Modes allows the vehicle spun-side to remain spinning and continues to utilize the GNC reaction wheels, as opposed to switching to thruster control. The other Safe Mode powers off the spin electronics to spin down, enters a contingency communications state, and switches to thruster control for GNC. Both modes enforce a GNC submode that maintains a sun-pointed rotisserie attitude (head-over-heels) rather than the science attitude.

Determining which SFP responses to map to the two Safe Modes was an important part of the FP design. Mapping to the Safe Mode that spins down the vehicle results in a significantly longer recovery timeline in operations, so it was utilized for a critical subset of faults or as part of a later tier in a multi-tiered response. The less severe Safe Mode that maintains vehicle spin was utilized for faults unrelated to the spin hardware or the reaction wheels that warranted a more substantial response than that enforced by the "Standby" mode. In the end, the decision to have multiple options for enforcing a safe state was valuable, as it allowed for different levels of SFP response severity that were more specific to the fault type. It also allows for operational timeline considerations, and for future SFP changes to utilize existing modes for re-mapping old responses or creating new ones.

CONCLUSION

During SMAP Commissioning, System Fault Protection software autonomously protected the vehicle from multiple faults to critical hardware, and the operations team successfully returned the observatory to its science state. The SFP also performed well in the presence of anomalous behavior below true safety limits by not taking unnecessary response actions and instead allowing the operations team time to monitor the behavior.

Many aspects of the SFP design can provide more visibility into anomalous behaviors. Having high-water mark data on SFP persistence counts aided in tracking and dispositioning of expected and unexpected behaviors. The team's use of ground alarms allowed for quick anomaly notification, and the fault protection dashboard provided a color-coded means of quickly assessing SFP state. Having telemetered values for all SFP error conditions would have greatly aided the FP team in tracking monitor threshold percentages, and an intermediate "yellow" monitor state would have also simplified tracking of anomalous behavior below the persistence limits.

In addition to high visibility, it is important that fault protection be easily modifiable during operations, with the caveat that critical and well-tested behaviors remain hard-coded to require more rigorous testing prior to modification. The decision to parameterize fault protection persistence and thresholds was very useful in ops as it allowed the team to tune the monitors based upon emerging behaviors. The use of global variables within monitor error checks and/or response actions can also provide the ability to modify new SFP, especially if the anomalous device behavior is not well understood. SMAP's use of a system modes state machine provided additional flexibility post-deployment and spin-up, as the modes could be easily repurposed for new operations scenarios. Finally, sufficiently staffing the fault protection team during Commissioning and early operations was essential to support anomaly recovery, SFP tuning, and flight software updates.

When multiple unexpected peripheral device resets occurred in Commissioning, the team faced challenges in implementing new SFP behaviors to recover the device locally prior to safing. Access to an Engineering Model would have provided valuable data on device reset behavior, sufficient test venue fidelity, particularly for complex devices, is highly recommended. As a work-around, the team was able to make use of the tiered response architecture and global variables in order to implement a flexible response to future resets. Overall, it is important that SFP aim to protect hardware that is not critical to safe mode and to attempt recovery of hardware that is essential to safing – within the requirements of the mission.

Finally, having different "levels" of safing within the system modes state machine allowed for more fault-specific system level responses. It also allowed for future re-mapping of fault protection responses to the less-harsh "Standby" mode for a subset of faults where preserving and maintaining mission science was an equal priority.

ACRONYMS

BAPTA	Bearing And Power Transfer Assembly
CDH	Command & Data Handling
DP	Data Product
EVR	Event Report
EM	Engineering Model
FP	Fault Protection
FSW	Flight Software
GNC	Guidance, Navigation, and Control
HPA	High Power Amplifier
HWM	High Water Marks
IMU	Inertial Measurement Unit
LWM	Low Water Marks
RBA	Reflector Boom Assembly
SAR	Synthetic Aperture Radar
SFP	System Fault Protection
SMAP	Soil Moisture Active Passive
SRU	Stellar Reference Unit
V&V	Verification and Validation

REFERENCES

[1] *SMAP (Soil Moisture Active Passive) Handbook: Mapping Soil Moisture and Freeze/Thaw from Space*. 2014 July. http://smap.jpl.nasa.gov/mission/description/

[2] Meakin, Peter; Jacome, Raquel. *Soil Moisture Active Passive Mission: Fault Management Design Analyses.* AIAA Infotech@Aerospace 2013 Conference, Boston, MA, 19-22 Aug. 2013.

ACKNOWLEDGEMENTS

The authors thank the members of the SMAP Fault Protection team for their invaluable contributions to the SFP design, testing, delivery, and operations—namely Julie Wertz-Chen, Raquel Jacome, Danielle Marsh, Alex Moncada, Gene Lee, and Gorang Ghandi.

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contact with the

National Aeronautics and Space Administration.

BIOGRAPHIES

***Jessica Clark** received a B.S. and Masters Degree in Aerospace Engineering from Georgia Tech in 2010 and 2012 respectively. She has been a systems engineer in the Fault Protection and Autonomy group at the Jet Propulsion Laboratory since 2012. While at JPL, she has served as the Project Systems Engineer on a CubeSat project, RACE, as well as a member of the operations team for MSL. On the SMAP project, she worked as part of the System Fault Protection team starting at project CDR, and is currently the SFP operations lead for the science phase of the mission.*

***Peter Meakin** is the fault protection lead for the Soil Moisture Active Passive (SMAP) mission, and has worked on a variety of projects at JPL. He is currently the flight system behavior architect on the proposed Europa mission, and also supporting the Fault Protection Design and Testing on InSight. Peter received his BS in Mechanical and Aerospace Engineering from Cornell University and received a minor in Applied Mathematics. He has contributed to the fault protection design for SMAP, Cassini and Constellation as well as support for Mars Exploration Rover Pan Cam image calibration. In addition to his interest in fault protection he has worked in the area of attitude control systems (ACS) engineering, supporting the development of ACS architectures for over numerous mission concepts. He was the ACS engineer for JPLs Titan Saturn System Mission and Jupiter Europa Mission studies and supported SMAP and Cassini attitude control teams.*