A Taxonomy of Uplink Planning & Sequencing Architectures for 16 Successful Space Missions (2001 – 2013)

David Mohr
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Dr
Pasadena, Ca, USA

Contents

1.0 Abstract

With a wide range of outstanding spacecraft, NASA's Jet Propulsion Lab (JPL) has accomplished a tremendous amount in space exploration, delivering a wealth of knowledge to the world. Excellent operations personnel, along with a well-built uplink process, are absolutely essential for a space mission to deliver on its objectives. This paper performs a taxonomy of the uplink Planning & Sequencing (P&S) system architectures that JPL has developed and executed for 16 successful flight projects over the past twelve years (2001-2013). In no particular order, the flight projects are

Galileo	Stardust	Mars Odyssey	Mars Exploration Rovers (MER)
Deep Impact	Spitzer	Genesis	Cassini
Phoenix	JUNO	GRAIL	Mars Reconnaissance Orbiter (MRO)
EPOXI	NExT	Dawn	Mars Science Lab (MSL)

The resulting P&S categorization (shown in Figure 3.5) does not necessarily line up with a typical categorization of space missions, such as orbiters, flybys, landers, etc. Here, flight projects and their P&S systems have been classified based on criteria that include a) level of *operational interactivity of the payload*, b) whether operations are *predictive or reactive in nature*, and c) whether a mission's *science opportunities are recurring*.

There has long been an effort to create a single multi-mission P&S architecture that will accommodate any mission type. Nonetheless, different architectures are the reality. Commonalities are discussed, but it seems clear that there are aspects that are necessarily unique.

The goal of this writing is to describe how different flight projects perform P&S, and what drives the need for different P&S architectures. This effort also addresses the degree to which each of the above criteria can be controlled by management, and what a project's goals should be. There are decisions that a flight project can make, whether early in development or entering an extended-mission phase, that will reduce operational complexity, cost, risk, and command errors.

2.0 Overview of Planning & Sequencing (P&S) Operations

Uplink operations personnel are on the front lines of accomplishing space mission objectives and, at times, it can feel as though the whole world is watching – especially during mission critical events such as orbit insertion or entry, descent, and landing. In fact, the media is very often present to see operations personnel hugging, laughing, and high-fiving in jubilation, or occasionally hanging their heads in sorrow. Safely and successfully commanding a spacecraft is of utmost importance in space exploration.

In a very real sense, P&S personnel are system engineers – requiring an understanding of the science/mission objectives, the spacecraft's instrument and engineering subsystems, and the environment in which the spacecraft is operating. It's the responsibility of operations personnel to command and utilize the spacecraft to safely collect the science data that will realize a mission's overall goals.

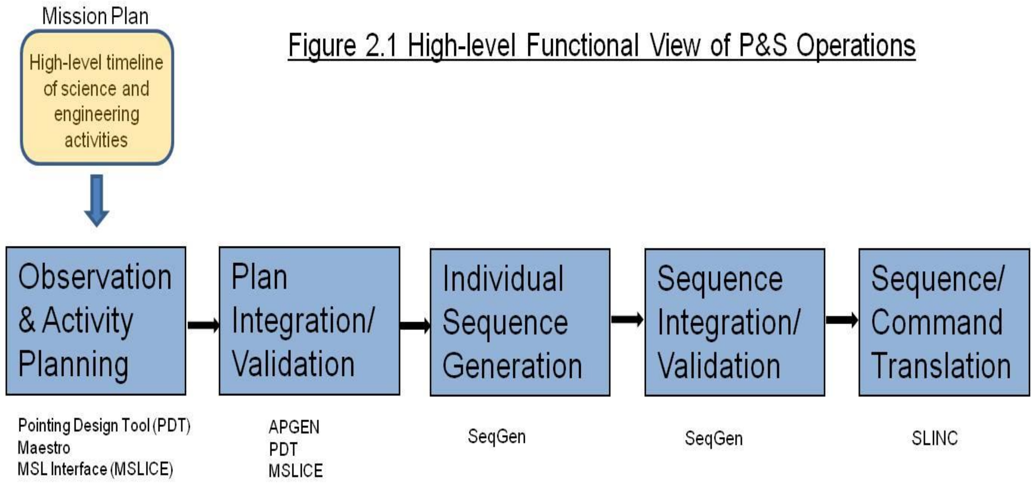

Figure 2.1 High-level Functional View of P&S Operations

Figure 2.1 shows the high-level process of P&S operations, along with common software tools that are used. P&S involves planning the science and spacecraft activities, generating command sequences that execute those plans, integrating those sequences, validating them, converting them into the correct uplink format, and then approving them for transmission to the spacecraft. (As an aside, ask anyone who's ever had the formidable duty of signing their name to approve a command load that will then be sent to a spacecraft if they recall the first time they did it – that's a guaranteed "Yes"!).

The sequence validation process typically uncovers a number of violations and conflicts. Some of these can be resolved by a single team editing their sequence product, and some resolutions will require coordinated changes across multiple teams. The integrated product is further analyzed by the spacecraft engineering team against their own spacecraft models, and by manual review.

2.1 Cassini Example

The complexity of P&S operations can vary greatly from project to project. For example, creating and validating plans and sequences for Cassini is a significant systems engineering effort. Consider the following scenario in which Cassini, currently in Saturn orbit, has an upcoming low-altitude flyby of the high-value target, Titan. After significant science planning effort to negotiate pointing (spacecraft attitude), data volume, and power modes, the following high-level plan (Figure 2.2) has been worked out –

Time Relative to Titan closest approach (C/A)	Main Activity	Details
-10hrs to -7 hrs	Thermal IR Observation of Titan Limb	Other remote sensing instruments also collecting data
-7hrs to -3hrs	Imaging various Titan regions (mosaic)	Other remote sensing and in-situ instruments also collecting data
-3hrs to -1hr	RADAR (Radiometry)	In-situ instruments also collecting data
-1hr to -40min	Transition from Reaction Wheels to Thrusters	Must be on thrusters when below 1300km Titan altitude. Enginering Team in control.
-40min to +1hr (incl closest approach)	RADAR (SAR)	Detailed surface imaging; some in-situ instruments (esp. Mass Spectrometer) also collecting data; some inst. must be commanded to sleep for power purposes.
+1hr to +1:20hr	Transition from Thrusters to Reaction Wheels	Engineering Team in control during transition.
+1:20hr to +3hrs	Near IR measurements	Other remote sensing and in-situ instruments also collecting data
+3hrs to +8hrs	Thermal IR measurements	Other remote sensing and in-situ instruments also collecting data
+8hrs to +18hrs	Downlink to Earth	Must have 70m DSN coverage in order to empty the SSRs (~3.5Gb)

Figure 2.2 High-Level Plan for Cassini Titan Flyby

If all goes well, the execution of this plan could acquire some ground breaking science data that will enhance the world's understanding of Titan's complex atmosphere and exotic surface features, and perhaps further our understanding of terrestrial bodies in general.

However, this plan is far from being able to be uplinked to the spacecraft. As is common for this type of plan, when it is sequenced and integrated, P&S software and other analysis uncover a number of problems. Some of the key issues with this particular plan (Figure 2.3) are the following

Errors/Conflicts	Message	After further analysis	Resolution
Turn interrupts	Software reports that a new turn is being commanded before a previous turn has completed (near C/A - 3hrs).	The Imaging Team's mosaic took longer to complete than planned. This is interfering with the start of the RADAR observation.	Rework of the Imaging Teams observation. Must then be resubmitted, and remerged.
IR Radiator Heating (>5 deg) Flight Rule violation	Software reports that the 5deg radiator heating threshold was crossed at +1:33.	Radiator heating began prior to C/A, but the 5deg threshold was not crossed until 13min into the IR observation. This heating can be caused by a low angle to the Sun, Saturn, the rings, or Titan.	Multiple changes involving RADAR, Engineering , and IR teams must occur.
Optical Remote Sensing (ORS) focal-plane-to-Ram angle < 80deg	Software analysis shows the "ORS boresight to Ram" angle goes below 80deg at C/A-00:33.	At this point the spacecraft is in Titan's atmosphere; particles can damage focal planes if pointed near the spacecraft velocity direction.	RADAR Team must rework observation design.
IR focal plane within 15deg of Sun	Software reports the "IR boresight to Sun" angle goes below 15deg at +1:24	As the IR Team begins the turn to their desired attitude, they inadvertently point too close to the Sun.	IR team will have to rework their observation design, likely adding turn time and reducing valuable observation time near C/A.
One of the reaction wheels (RWAs) enters a low-RPM region	Software reports RWA#2 is below 300RPMs from -4:10 to -3:05	Extended periods at low RPMs causes excessive friction. Something in the pointing profile must be changed.	New secondary pointing must be negotiated with the Imaging Team and others collecting data at this time.

Figure 2.3 Subset of Error Report for Titan Flyby Period

Most of the issues listed in Figure 2.3 do not appear in software when planning an individual science observation. They can only be realized when analyzing the merged plan/sequence. These issues typically need to be worked with multiple instrument teams and engineering subsystems to find resolutions. Also, this is but one day of a 40-day sequence so the full list of "Errors/Conflicts" is significantly larger. As a note, once all of the issues have been cleared up and the sequence validated, the spacecraft must execute the command load at a great distance from Earth – in Cassini's case, 1 billion km and a one-way light time of ~90min.

This is a typical P&S scenario for Cassini, where spacecraft activities, and in turn, sequence development, are highly dependent and integrated. However, with other missions there is a wide range in the level of interface between the science instrument teams, the engineering team, and P&S personnel. For the Mars Exploration Rovers (MER) and Mars Science Lab (MSL), the science instrument teams have significant interaction early in the planning cycle (through Science Planning), but then minimal after their individual sequences are delivered to the Sequence Team. Still with other missions, the instrument teams' interactions are further reduced – e.g., for Odyssey and Spitzer it is the norm for an instrument team to directly send their non-interactive commands (and on-board block calls) to the spacecraft, without any formal project approval. Reasons behind these differences are explained below.

3.0 What are the Categories?

There are many different ways that flight projects can be categorized. Perhaps the most obvious is simply to say there are flybys, orbiters, landers, astronomy (telescope), sample returns, and touring (e.g., Cassini) types of missions. However, in attempting to categorize P&S systems, they do not necessarily line up with these *mission* categories. There are even a number of ways that *P&S systems* may be classified; here's one way that could prove useful –

Categorize them based upon the following:

1) Are the operations *predictive* or *reactive* in nature?

2) To what degree must the science instruments share and compete for spacecraft resources to accomplish their objectives (i.e., is the planning *interactive or non-interactive?).*

3) Is the mission generally characterized by science opportunities that are unique/rare, or are they recurring?

Here we answer, "*What combination of these criteria does each P&S system exhibit?*" and "*What is meant by each of these criteria, and why were they chosen?* "

3.1 Are the operations *predictive* or *reactive* in nature?

This question can typically be reduced to – does the flight system have significant interaction with its environment – driving, drilling, digging, etc.? Even more specifically – does the current plan in development depend on outcomes from the previously executed plan (as described by the latest telemetry)? If so, then the P&S system is *reactive*, in that it must *react* to events such as new science opportunities, a change in the predicted available power, or the vehicle being in a different location or orientation than was planned (e.g., as a rover might). Figure 3.1 shows the architecture for a reactive P&S system.

Figure 3.1 Reactive Planning & Sequencing
(e.g., MER, Phoenix, MSL)

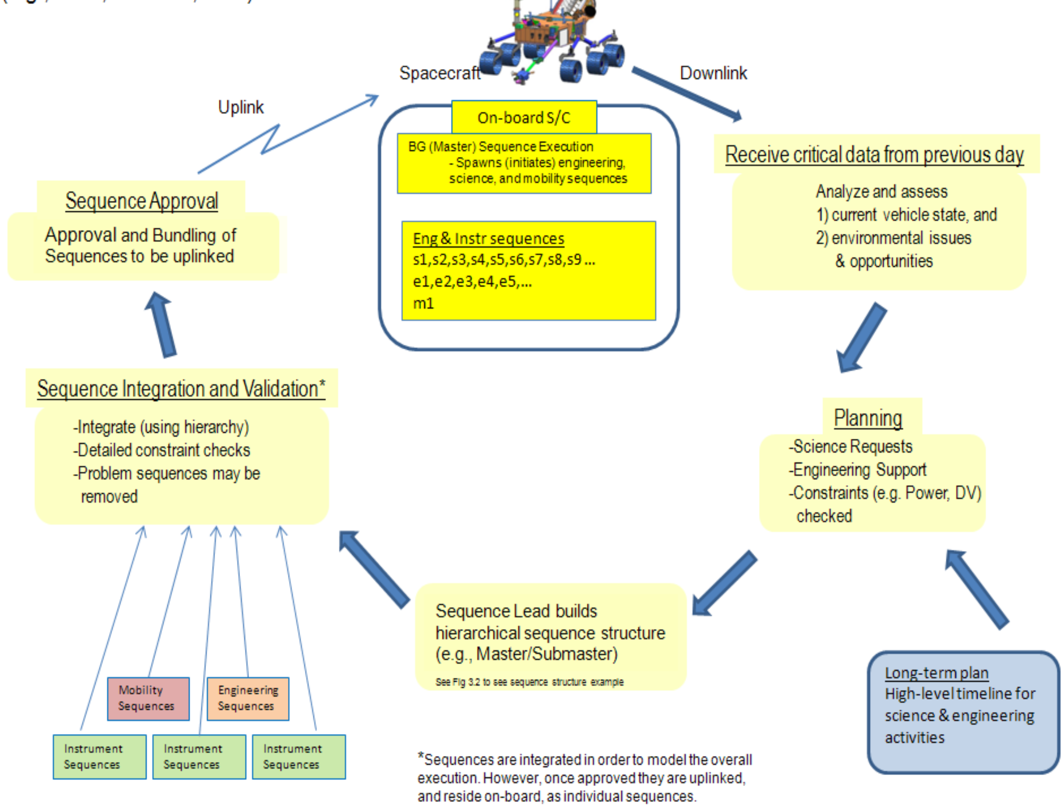

As a general statement, currently only operations of landers (including rovers) are described as "reactive", and because of the unpredictability involved, P&S occurs on a very short timeline. For example, a Mars Exploration Rover (MER) plan and sequence are developed in a single day, and the sequence nominally executes in a single Martian day (or Sol). The unpredictability also explains the need for a hierarchical "Master/Submaster" sequence structure. A master sequence is needed to declare exactly when the overall plan will begin and end, when the vehicle CPU will turn on and shut down, and when communication windows occur, as well as other engineering activities. Master sequences call submaster sequences. Submaster sequences can have logic built in that will have uncertain execution times and durations, hence the need for a master to have absolute timing control. Drives and arm sequences are typically submasters, or can be called by submasters. Instrument sequences can be run or initiated from within drive and arm sequences, and often are. "Run" is used when we want the *calling* sequence to wait until the *called* sequence has completed execution before continuing, and "initiated" is used when we want the *calling* sequence to continue its execution in parallel with the *called* sequence. Figure 3.2 shows an example of a Master/Submaster sequence structure.

Figure 3.2 Hierarchical (Master/Submaster) Sequence Structure
For one day (sol)

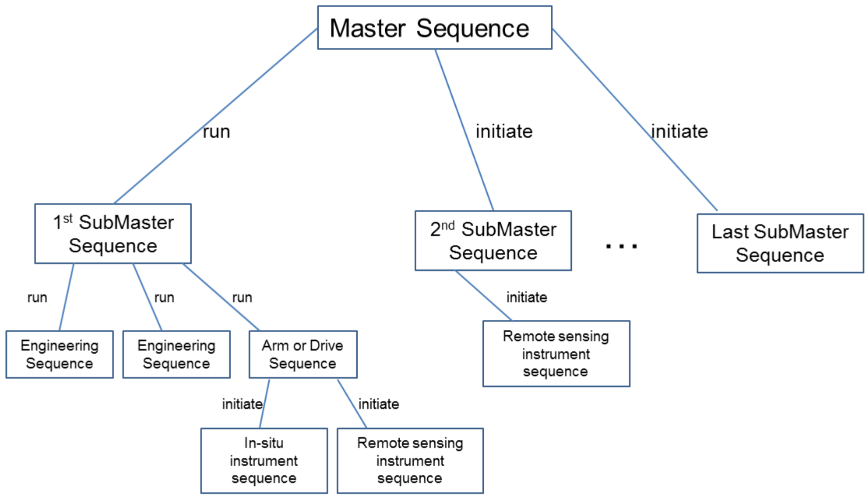

MER, Phoenix, and MSL all fall squarely into this "Reactive Ops" category. For these missions (and future missions requiring Reactive P&S) there is, and should be, a strong desire to reduce payload interactivity. There simply is not enough time to address competing science priorities and desires in real time during a reactive operations planning day. Interactivity also increases the likelihood of conflicts that must then be resolved in this brief time period. For reactive operations, a multi-week strategic plan is typically in place that lays out the general activities for each planning day. Analyzing the results of the previous days' activities, creating and approving the new detailed activity plan, building the required master/submaster sequence structure, then creating, merging, and validating the sequence products, is all difficult enough on a short timeline without adding in science contention and negotiation.

Predictive operations refers to those missions where the spacecraft trajectory is known well ahead of plan development. With the ephemeris of the spacecraft and the target(s) (e.g. planet, moon, asteroid) understood, science opportunities are known with large lead times. This permits science observations, engineering activities, and the integrated plans and sequences to be developed well ahead of the time they will need to execute on-board the spacecraft.

For example, Cassini performs many orbit trim maneuvers (three per orbit), to keep the spacecraft on its planned trajectory. This means that the timing and geometry will be well known for each periapse, satellite flyby, ring plane crossing, and solar/stellar/Earth occultation, as well as the associated target info such as solar lighting angles and spacecraft-target relative motion.

3.2 Will operations involve activities that are generally *interactive or non-interactive?*

Can the instrument teams perform their science activities without impacting other instrument activities? What spacecraft resources must be shared and competed for? – power, pointing (spacecraft attitude), data volume, data collection rate?

A deep space telescope mission (e.g., Spitzer) or missions with a circular, repeating orbit, such as Mars Odyssey, will typically have their science instruments pointed in their preferred direction for data collection during nominal operations. Also, missions such as these often have sufficient power and data volume capacity for all of the instruments to collect data simultaneously without impacting each other, or the engineering subsystems. Furthermore, if the high-gain antenna (HGA) is gimbaled such that the spacecraft can downlink data to and receive commands from Earth without interrupting science operations, the system is perhaps as non-interactive as it gets. A P&S system like this will require a relatively low-level of operations support.

MRO, which has several remote sensing instruments, has a pointing profile that is somewhat more interactive. Science Planning captures science target requests (i.e., latitude, longitude points on the Martian surface) for the instrument teams, and produces an integrated target list. These non-nadir observations require the spacecraft to roll to attitudes that are often not useful to other instrument teams. The Deep Impact (DI) flyby spacecraft, with its three science instruments sharing multiple resources, certainly had some level of interactivity as well. The DI flyby sequence took several years to develop, re-negotiate, and test before it was finalized.

On missions such as Cassini, power, pointing, data volume, and data collection rate are *all* shared by a large instrument suite, and negotiations occur for each resource. Highly complex, interactive plans and sequences must be developed in order for each instrument to obtain its required data. Not only does high interactivity mean that spacecraft resources must be competed and negotiated, it also typically means there will be conflicts when the plan/sequences are integrated. Hence, several merges and iterations are needed to get a valid, conflict-free sequence. Figure 3.3 shows the architecture for a predictive, interactive P&S system.

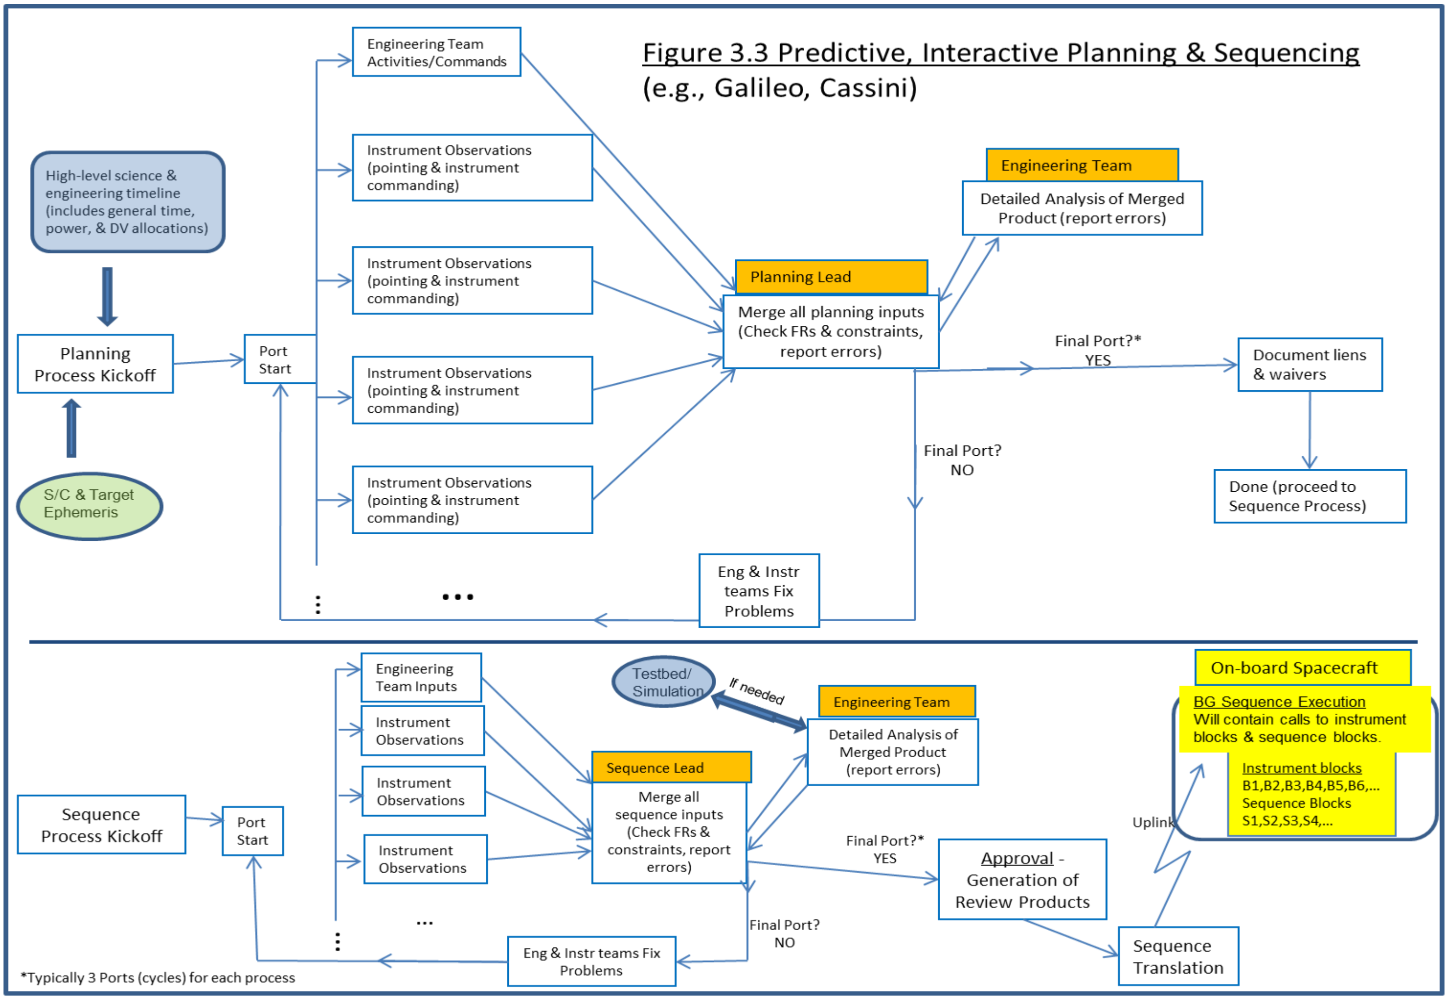

Figure 3.3 Predictive, Interactive Planning & Sequencing (e.g., Galileo, Cassini)

Complex planning leads to complex sequencing, and the need for several cycles to get it right. This directly leads to a more lengthy P&S process. When the length of the process exceeds the length of sequence execution, the project will then be required to develop more than one sequence simultaneously. P&S costs have a direct correlation to complexity and payload interactivity.

In some cases, an activity performed by one instrument can completely preclude several other instruments from capturing data at all. Radio Science and RADAR activities often use such a high level of power that other instruments must be put in sleep mode during that time.

The MER operations team maintains a longer range plan that declares a baseline “sol type” for each particular sol. This means the general types of activities are known prior to the start of planning – is it a *drive* day, an *arm* day, or are we going to focus on *remote sensing* today? Even though each P&S day starts with a general plan, there are negotiations and decisions to make on a short timeline. The first several months of MER operations were especially contentious and included quite a bit of interaction. Remote sensing instruments on the mast often had observation requests that directly impacted each other. Power issues were also substantial early on, and the planning process typically included the removal of several science requests.

The introduction of a "skeleton plan" template and data volume recommendations, to go along with the long-term plan, significantly constrained (and simplified) the science planning effort. Furthermore there is typically no urgency to perform a particular science activity on any given sol; the science target will very likely be available on the following sol.

MER has successfully honed its P&S system to be generally *non-interactive*, thus lessening the need for complex planning and science negotiation. Over time, the flight team has developed and maintained a matrix of incompatible activities that guides daily plan development. Still, power and data volume must be estimated for each activity and modeled for the entire plan; this often requires editing and rework during tactical planning days.

With twice as many science instruments on MSL than MER, as well as highly-complex arm activities, there is much greater potential for payload interactivity as teams contend for the limited spacecraft resources. For this reason, MSL developed strategic processes for pre-approval of plans, as well as various activities, prior to the start of a tactical planning day. Multi-mission projects such as Mars Odyssey and Dawn have also been successful at reducing operational interactivity. Well-defined mission phases give priority to specific instruments, essentially eliminating any contention for pointing, data volume, and power. Multi-mission flight projects are those that utilize a standard set of planning, sequencing, and command generation processes to meet its objectives. This is done whenever the mission design and operations concept allows for it.

As a general rule, if interactivity can be reduced or eliminated during operations, then it definitely should be. The more interactivity within the science payload, the more complex and expensive the P&S system will be.

Sometimes there are reasons why some level of interactivity is unavoidable, and even beneficial; for example,

i) If a mission has a significant variety of science targets (e.g., surface, atmosphere, satellites, magnetosphere, rings), then the sheer number of science goals could be very ambitious, requiring a considerable number of instruments. This will likely increase interactivity.

ii) Data from one instrument can provide scientific insight or physical context to other instruments (e.g., a spectrometer can utilize a RADAR or camera image to help gain perspective on what they were viewing when they acquired certain data).

Cassini demonstrated real synergy when they employed a significant portion of their 12-instrument suite to discover liquid water reservoirs and geysers on Enceladus. The magnetometer noticed a strange draping of Saturn's magnetic field near Enceladus' south pole consistent with a local atmosphere. Ultraviolet observations of a star being occulted by Enceladus' south polar region showed a signal attenuation that was also consistent with a local atmosphere. Then, visible light images at high phase angle gave visual evidence of the liquid plumes. Subsequent mass spectrometer (and UV spectral) data revealed that these liquid plumes were indeed more than 90% H_2O.

While there may be scientific benefits to certain interactive science operations, interactivity inevitably leads to plans and sequences of higher complexity. This, in turn, lends itself to higher costs, increased error rates, and more risk. Interactivity is perhaps the biggest driver of P&S system cost. This is true in terms of both software effort and operations personnel. A Mission System manager should strive to reduce interactivity as much as possible early on.

3.2.1 Multi-Mission Projects and Non-Interactive Commanding

It is not coincidental that essentially all of the missions that use multi-mission P&S fall into the "Non-Interactive" category. In the multi-mission world, the term "non-interactive" is specifically used to describe the majority of payload commands and files that are sent to the spacecraft; they are unable to impact any other instrument activities, or adversely affect resource margins or spacecraft health & safety.

Non-interactive commanding provides instrument teams the ability to command their instrument without interacting with the rest of the operations team, and without going through a formal approval process. Although not always possible, to achieve the simplest, lowest-cost, lowest risk P&S system, this non-interactive multi-mission scenario is what a project should strive for. Figure 3.4 shows the Multi-mission P&S architecture which is generally both predictive and non-interactive.

Some flight projects can even be described as *extremely non-interactive*. For these, nominal operations is such that *all* of the science commanding is sent directly to the spacecraft, without being merged and validated with the background sequence.

Figure 3.4 Predictive, Non-interactive (Multi-Mission) P&S
(e.g. Mars Odyssey, Spitzer, Deep Impact, Dawn, JUNO)

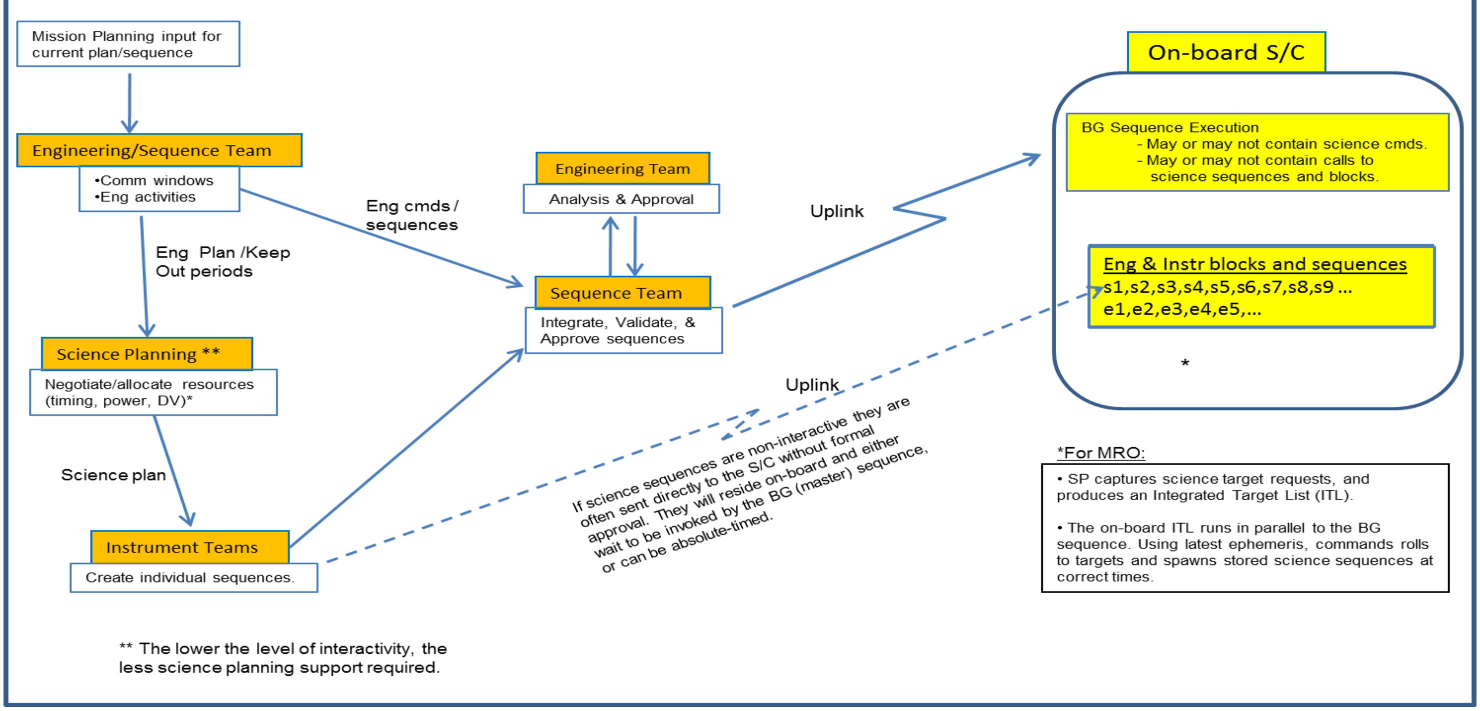

3.3 How many unique science opportunities will the mission have?

As posed here, a unique science opportunity refers to the relative location, or geometry, of the spacecraft and a science target that is rare or unique, and will not occur again for a long time, if ever.

Flyby missions, such as EPOXI, have a single unique science opportunity – the flyby itself. Flyby missions use multi-mission processes, and can be categorized as "Predictive, Non-interactive, *Recurring opportunities*", with the obvious exception of the one highly critical, unique science opportunity. This one sequence goes through many iterations and gets validated many times. Cassini has many rare science opportunities, such as satellite flybys and solar occultations, and some unique science opportunities as well, such as the Phoebe encounter and high-inclination ring views at Equinox. Some opportunities are unique because of the geometry involved, such as an ideal solar lighting on a Titan lake from the current view of the spacecraft.

By this definition, Rover missions do not have unique science opportunities, they are recurring. If a rover encounters an object of interest, the operations team can choose to stay at this location until the science teams have exhausted their desires. Mapping missions, such as Mars Reconnaissance Orbiter (MRO) and Mars Odyssey, also do not generally have unique science opportunities (as defined above) – the science opportunities are *recurring*. This certainly does not mean that mapping missions and rover missions do not make ground-breaking science discoveries, because they absolutely do. It simply means that the time window they have to perform their high-value science observations is much larger, allowing the entire payload to participate.

Having a large number of unique science opportunities *and* heavy payload interaction can be a troublesome combination requiring a significant level of science planning support. On Cassini, for example, power, data volume, spacecraft pointing, and data collection rate are all shared, negotiated resources. Furthermore, because the RADAR and the remote sensing instruments are oriented 90^{o} apart on the spacecraft, they generally cannot observe simultaneously. If an upcoming science opportunity is rare (such as a low-altitude Titan encounter), there can be significant contention as to which team will be allocated the closest approach period; contention which must be settled. Going to Saturn, in itself, is a rare opportunity. Combine that with a complex "touring" trajectory where each orbit has different science opportunities than every other, and the effort to produce and validate plans that satisfy science team goals can be sizable. The cost for this type of P&S support will be high.

Unique science opportunities are certainly not a bad thing. But when predictive operations (i.e., significant lead time) is combined with a rare/unique opportunity, there will be a desire to "maximize science". From there, a highly complex, interactive plan and sequence must be developed for each instrument to acquire its data.

Figure 3.5 shows a summary of the P&S architecture categorization listed by mission.

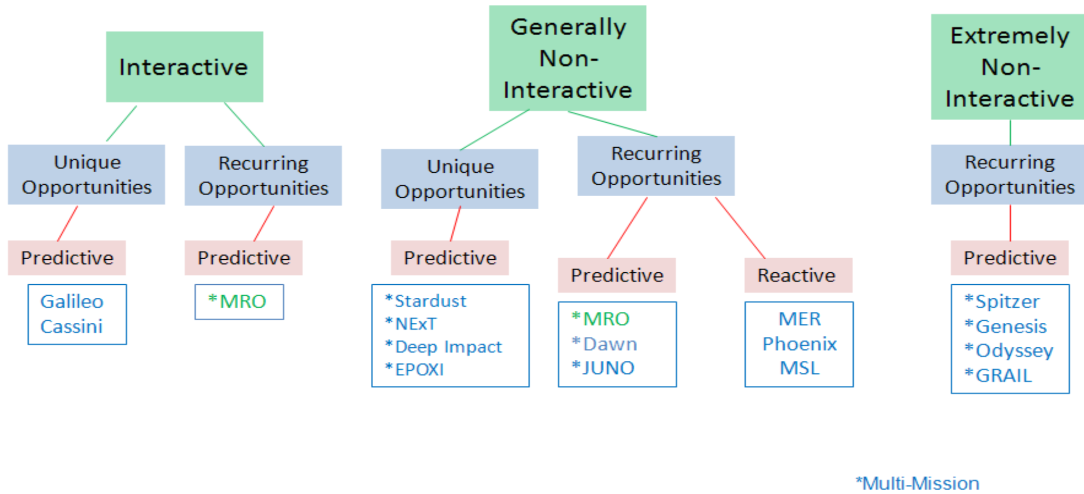

4.0 <u>What aspects can management control? … and what should be the goal?</u>

4.1 <u>Reactive or Predictive?</u>

Managers may have little to no control as to whether their nominal operations are predictive or reactive. This will be dictated by the mission concept – orbiter, flyby, lander, etc. Often these environments are very harsh, bringing uncertainty to the flight system life expectancy. Hence, there is generally a real push to *utilize the asset while we still have it*. Currently, reactive operations provides the best way to do just that. It attempts to make best use of the most current state of the spacecraft to investigate the latest findings of the spacecraft's environment.

On the down side, performing the entire P&S process and commanding on a daily basis can be taxing on operations personnel, particularly if they are asked to work according to the spacecraft's local time (e.g., Mars time). This will drive the need for additional staffing. Also, the shear frequency of commanding will inherently increase the risk of command errors.

For reactive operations on MSL, multi-day plans/sequences can be accomplished if detailed spacecraft status is not needed after execution of the first day's activities prior to execution of subsequent days' activities. For example, driving on the first day, and then performing remote sensing observations on the next two. Additional power margin should be kept for activities after the first day in case of unexpected power decreases. On-board automation, such as MER's use of AutoNav, will continue to reduce the reactive operations support needed on the ground. Also, there is an evolving capability for on-board software to perform its own assessment as to spacecraft location, orientation, safety, and authorization to perform certain activities. In addition, there is also a growing ability for spacecraft to locate and assess science opportunities autonomously.

4.2 <u>Interactive or Non-interactive?</u>

Management's ability to control payload interactivity can be limited as well, especially if spacecraft operability is not sufficiently considered in a project's early phases. If instrument teams cannot perform the activities they need without impacting each other as far as pointing, power, and data volume, then there may be little that can be done. There *will* be a significant science planning presence needed to negotiate these resources and to produce an optimized plan. Elaborate software will also be needed to model and work out the detailed plan for the shared use of these resources. Additionally, there will inevitably be conflicts to work out as this complex plan/sequence is iterated toward the clean, final, uplink product. This quickly gets expensive!

Spacecraft pointing can be the most interactive resource of all. Different science instruments study different things – some want to observe ground features, some observe atmosphere (and at various altitudes), some observe magnetospheres, some only plasma, etc. If these instruments are not able to articulate individually, then the entire spacecraft must be articulated and pointed for instruments to capture their data. If changes in spacecraft orientation must be performed to collect science data, or to downlink, then *pointing exclusion zones* will likely exist. Solar arrays must stay *within* a certain plane or angle to the Sun, and radiators and remote sensing focal planes typically must stay *outside* of certain angles to the Sun (and other bright bodies). Complex software for pointing, tracking, data volume, power and thermal modeling will be needed for science teams planning their individual activities. When science teams add complex observation design requirements, such as scanning and mosaicking techniques, the expense can grow significantly in terms of planning software, instrument team time and effort in creating these observation designs, Attitude Control (ACS) and other Engineering team analysis and constraint checking, and in taxing the spacecraft ACS system (e.g., degradation of reaction wheels). *Just how* the spacecraft will need to operate in order to accomplish mission and science objectives should be carefully considered early in flight system development.

Increased interactivity and complexity generally means additional merge cycles and longer development times will be required to develop constraint-free sequences. When P&S development time exceeds sequence execution time, the project will have to develop multiple plans and sequences simultaneously; this likely means *even more* staffing support will be required.

Hence, to the extent that a newly forming project can minimize interactivity and still meet mission objectives, it *should*. In general, when trying to reduce interactivity, think from the perspective of each instrument team, "What's standing in our way when we want to capture our required data? Are there data volume or power limitations? Are we not able to point when/where we want?"

4.2.1 Reduction of Operational Interactivity

A developing flight project may want to consider the following when attempting to reduce operational interactivity:

- Size the spacecraft power source to have sufficient power for nominal operation of the engineering subsystems and the entire instrument payload, as several missions have done.

- Eliminate data volume contention prior to launch by allocating each instrument a pre-sized hard partition of the on-board recorder for its raw data. This was successfully done by MRO; even if the instrument teams *wanted* to negotiate data volume during operations, they can't.

- Clear and thorough operations scenarios are critical during the very early phase (e.g., proposal) of a project, as it can impact flight system design. If possible, have instruments nominally oriented in their preferred direction of data collection (e.g., mass spectrometer to ram direction, remote sensing to nadir) - allocate resources for a scan platform if needed. Also, consider options regarding individual instrument articulation when warranted. An articulating HGA will also provide major reduction in pointing interactivity.

- Co-align cameras and other remote sensing instruments. Most missions already do this, but different remote sensing instruments will have different observation requirements – some prefer ‘staring’ at a target for various durations, some want to mosaic, some want to scan, etc. – each impacting the other.

- If there *are* instruments with conflicting observation requirements (e.g., two instruments requiring significant power or different spacecraft attitudes), then create operational scenarios that establish priority periods. Allocate blocks of observing time where each instrument is given a turn as the highest priority.

Although less apparent and probably more disruptive, there are steps that a project entering extended-mission phase can take to reduce interactivity and complexity. Once primary science/mission goals have been met, a project may be forced to perform P&S with a significantly reduced budget. If this occurs, a project may want to consider the example below. The result *will* be a reduction in science return, but that is arguably better than trying to maintain the same level of interactivity on a reduced budget, and increasing operational risk.

As with projects in early phases trying to reduce interactivity, consider allocating blocks of observing time where a single instrument is given highest priority. The larger these blocks of time, the less interactive the plan and sequence will be.

Example: The Cassini Titan flyby plan shown in Figure 2.2 is a good example of a complex, interactive plan. And, as commonly occurs, several errors and conflicts occur when the sequence is merged. After 3 to 5 cycles of iterations and merges, this sequence of activities will be clean, and ready to be uplinked. In turn, valuable *multi-disciplinary* science information will result, but at what cost? Significant P&S support was required to make this happen.

If the team that is allocated the closest approach time (RADAR in this case) was given priority for the entire period, some remote sensing instrument data would be lost. However, 4 of the 5 “Errors/Conflicts” that appear in Figure 2.3 either would not have occurred or would have been flagged and fixed by the RADAR team (internally) prior to the first merge cycle. This new *lack of interactivity* leads directly to fewer errors, fewer cycles, less overall time to develop a clean plan/sequence, and less P&S required support.

Then perhaps a remote sensing instrument may be given priority throughout the next Titan flyby.

4.3 Unique or rare science opportunities

This only really presents an issue if there are some instruments that are unable to collect their data because of pointing or other resource limitations, or if a highly intricate plan must be devised for all of science teams to capture their desired data. So, again it comes back to interactivity.

If rare or unique opportunities are an inherent part of a mission, it is important to have clear operational scenarios as to how all instruments will collect their required data. Otherwise the mission will be saddled with contention and complexity – both of which are expensive.

5.0 Commonality in Planning and Sequencing

5.1 Commonality in P&S Software

P&S software is critical to the control of, and communication with, a spacecraft. P&S software is used for modeling, visualizing, planning, and development of science and engineering activities and sequences. It's also used for checking flight rules ad other constraints, and translating the validated command sequences for uplink.

Currently, and for the past several years, P&S software has been developed as two separate components. The multi-mission (core) software provides the general capabilities (e.g., planning and scheduling activities, packetizing commands). The core software is then *adapted* to meet the needs of a specific mission. The adapted part of the software involves modeling the activities needed for planning, converting the mission's command dictionary into models, the modeling needed for flight rule and mission rule checking, and developing blocks for reusable science and engineering activities.

By separating the P&S software components and development in this way, overall cost has been reduced. Verification of the core software element only needs to be done once. Each project can then just focus on verification of their adapted part of the software. [1]

5.2 On-board Commonality – Virtual Machines and Sequencing

All 16 flight projects mentioned in this paper utilize aspects of the concept of *virtual machines*. Flight software is a computer program running on a spacecraft's hardware. Part of this program, the sequence engine, is an environment that stores and executes commands and sequences. This environment is thought of as a virtual machine, with command sequences thought of as programs that run on these machines. This virtual machine has a simulated processor, memory, and registers.

Virtual machines support multi-tasking, as multiple sequences can be active at one time – each running on its own sequence engine. Sequences can contain programming logic, they can be event-driven, or they can simply contain lines of commands in time order. Thus, sequencing can be done on a conditional basis – i.e., wait for a certain condition before executing, as opposed to waiting the maximum time for an activity to complete.

Reusable blocks and sequences can be stored and invoked in these virtual machines. Blocks can have parameters passed to them from the background (master) or other sequence. Sequences can be run which poll the flight software for events or conditions passed as global variables. Thus, conditional statements can flag violated assumptions, and looping constructs can retry an action until it succeeds.

5.3 Multi-Mission Projects

Most JPL missions utilize a common P&S architecture, a common set of processes, and at times, can even interchange a common set of P&S personnel. By standardizing the core P&S operations, multi-mission projects have been successful at reducing development and operations costs, and meeting mission objectives with relatively small P&S teams. These missions can be characterized as non-interactive, with few rare/unique science opportunities. They are purposefully planned this way to reduce operational complexity and cost.

5.4 Why Don't All Missions Use Multi-Mission P&S?

In a sense, all missions *do* use a common P&S architecture. Figure 5.1 shows a generic view of P&S operations.

Figure 5.1 Generic View of Planning & Sequencing Operations

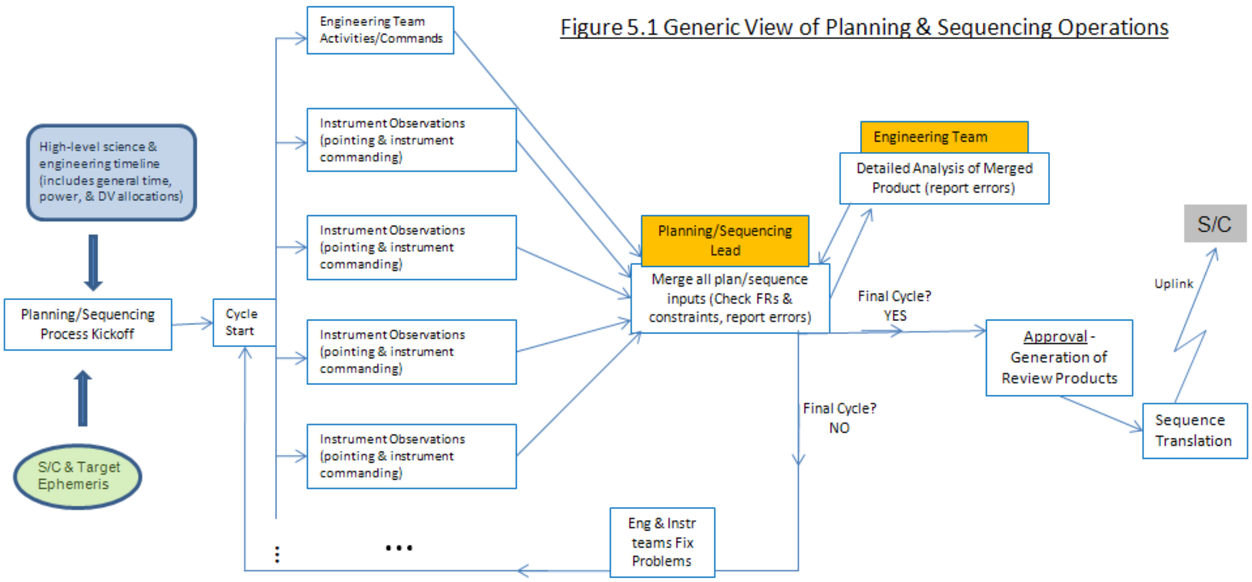

Functionally, it may be used to describe *any* P&S system. It's in the time and effort required to perform these functions where flight projects differ. Some of these boxes may represent a sizable effort for one project, but then may be omitted altogether for another. Important differences can be found in the following –

1) The level of Science Planning support required. For a highly interactive P&S system the science planning effort will be significant and just producing a clean *plan* can take several weeks. For a non-interactive P&S system the science planning effort is drastically reduced, and in some cases non-existent.

2) The number of cycles needed to produce the final sequence product. Highly interactive P&S systems nominally require multiple cycles to produce a constraint-free uplink product. For reactive systems and certain multi-mission systems, a single cycle (accompanied by occasional plan/sequence modifications) is typically sufficient. Still, for certain non-interactive multi-mission systems, command loads are often sent directly to the spacecraft without any formal approval (i.e., no cycle at all!).

3) The length of time to complete the P&S process. For a reactive system it may be a single day; for a multi-mission system it's typically three weeks; and for a highly interactive system it may be a few months.

6.0 Summary

Flight projects and their P&S systems can be classified based on a) level of *operational interactivity of the payload*, b) whether operations are *predictive or reactive in nature*, and c) whether a mission's *science* *opportunities are recurring*. This is what was used to categorize the 16 successful flight projects named in this write-up. Commonalities are discussed, but it seems clear that there are aspects that are necessarily unique.

Payload interactivity is the biggest contributor to why more missions have not been able to more fully utilize the “Multi-mission genre”. And when a manager or system engineer is looking for ways to simplify their P&S operations, their focus should be on reducing this interactivity. If interactivity cannot be controlled, then there *will* be significant P&S support needed to negotiate the spacecraft resources, to model and work out the detailed plan, and to clean up the inevitable conflicts as this complex, interactive plan is sequenced and iterated toward a conflict-free uplink product.

References/Acknowledgement

[1] L. Needles, “Cost Reduction from Multi-Mission Sequencing Software”, July 2003

The work described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

The author would like to thank the following individuals for their contributions to this effort: Dr. Kari Magee, Steve Wissler, Matt Keuneke, and Bill Heventhal.