Handling Late Changes to Titan Science

Jo Eliza Pitesky[1], Kim Steadman[2], Trina Ray[3] and Marcia Burton[4]
Jet Propulsion Laboratory, California Institute of Technology , Pasadena, CA, 91109

The Cassini mission has been in orbit for eight years, returning a wealth of scientific data from Titan and the Saturnian system. The mission, a cooperative undertaking between NASA, ESA and ASI, is currently in its second extension of the prime mission. The Cassini Solstice Mission (CSM) extends the mission's lifetime until Saturn's northern summer solstice in 2017. The Titan Orbital Science Team (TOST) has the task of integrating the science observations for all 56 targeted Titan flybys in the CSM. In order to balance Titan science across the entire set of flybys during the CSM, to optimize and influence the Titan flyby altitudes, and to decrease the future workload, TOST went through a "jumpstart" process before the start of the CSM. The "jumpstart" produced Master Timelines for each flyby, identifying prime science observations and allocating control of the spacecraft attitude to specific instrument teams. Three years after completing this long-range plan, TOST now faces a new challenge: incorporating changes into the Titan Science Plan without undoing the balance achieved during the jumpstart. Some changes add additional science opportunities on top of existing observations, as when we devised a new way to gather additional gravity data without impact to the originally planned science observations, using the spacecraft's Low Gain Antenna. Balance can also be impacted when instrument anomalies threaten the loss of a unique high-priority science opportunity. In response to one such situation, we created an alternative flyby timeline while keeping the original timeline viable late in the sequence planning process. As the aging spacecraft's capabilities change, we respond by tweaking long-planned observations. And as our consumables run low and project management scrutinizes their use ever more carefully, we add early detailed analysis of hydrazine use during Titan flybys, allowing us the option of redesigning (and thus saving) observations that might otherwise be removed during sequence development as being too "expensive". All this must be accomplished with a smaller workforce than was available during the Prime and first Extended mission. This paper looks at how TOST handles these and other late changes to Titan science

Nomenclature

CSM	=	Cassini Solstice Mission
HGA	=	High Gain Antenna
IDS	=	Interdisciplinary Scientist
LGA	=	Low Gain Antenna
MAPS	=	Magnetospheric and Plasma Science
ORS	=	Optical Remote Sensing
OST	=	Orbital Science Team
PSG	=	Project Science Group
RBOT	=	Reaction Wheel Assembly Bias Optimization Tool
RSS	=	Radio Science Subsystem
SCO	=	Spacecraft Operations
SEP	=	Sun-Earth-Probe
SOST	=	Satellites Orbital Science Team

[1] Science Planning Engineer, Cassini Project, 4800 Oak Grove Dr. M.S. 230-205, Nonmember.
[2] Science Planning Engineer, Cassini Project, 4800 Oak Grove Dr. M.S. 230-104, Nonmember.
[3] Deputy Science Planning Team Chief, Cassini Project, 4800 Oak Grove Dr. M.S. 230-205, Nonmember.
[4] Integration Scientist, Cassini Project, 4800 Oak Grove Dr. M.S. 169-506, Nonmember.

TOST = Titan Orbital Science Team
TWT = Target Working Team

I. The Mission, Spacecraft, and Instruments

The Cassini-Huygens mission to Saturn is a collaborative effort of NASA, ESA, and the Italian Space Agency[1]. The spacecraft launched on October 15, 1997 on a Titan IV-B/Centaur launch vehicle. After seven years, 3.2 billion kilometers (2 billion miles), and 4 gravity-assist flybys of other planets, it entered orbit on July 1, 2004. The spacecraft studied the planet, its rings, and its magnetosphere over the course of 76 varied orbits in the prime mission. To study Saturn's satellites, the spacecraft made targeted flybys of Phoebe, Hyperion, Dione, Rhea, and Iapetus, along with 3 flybys of Enceladus, and 45 of Titan. In summary, the Cassini prime mission was the most complicated gravity assist tour ever flown.[2] The Cassini Orbiter also carried along the Huygens probe, destined to measure Titan's atmosphere *in situ* and land on Titan's surface. The probe was deployed on December 25, 2004. Three weeks later, on January 14, 2005, it entered Titan's atmosphere and landed on the surface 2 hours later. The probe sent measurements and images to Cassini for transmission to Earth.

The spacecraft communicates with Earth largely through one high gain antenna but also carries two low gain antennas. Three radioisotope thermal electric generators provide power.

Cassini's twelve science instruments are grouped into three categories: Optical Remote Sensing, Fields/Particles/Waves, and Microwave Remote Sensing. The Optical Remote Sensing suite is comprised of a visible wavelength imaging camera (Imaging Science Subsystem, or ISS), an ultraviolet imaging spectrometer (UVIS), and infrared instruments (Cassini Infrared Spectrometer, or CIRS, and Visible and Infrared Mapping Spectrometer, or VIMS). The Fields/Particles/Waves suite is comprised of a magnetometer (MAG), cosmic dust analyzer (CDA), radio and plasma wave system (RPWS), ion and neutral mass spectrometer (INMS), plasma spectrometer (Cassini Plasma Spectrometer, or CAPS, the only instrument currently not working), and a magnetospheric imaging instrument (MIMI). The Microwave Remote Sensing suite is comprised of RADAR and the Radio Science Instrument (RSS), both of which use the high-gain antenna as an instrument.

Figure 1 identifies the science instruments. The Cassini mission requires operations on a global scale, and multiple time zones. In the final spacecraft configuration, the instruments were all mounted to the body of the spacecraft instead of a scan platform, which posed the single greatest challenge to operation complexity. The entire spacecraft must be rotated for any one instrument to achieve a desired target, and then the entire spacecraft must be rotated to point the high-gain antenna to earth to downlink the collected data. However, the optical remote sensing instruments are roughly co-aligned so they can often collect data collaboratively. On a typical Titan flyby the spacecraft collects science data for 30-40 hours by pointing the spacecraft at a variety of targets. One instrument at a time controls the pointing of the spacecraft, and other instruments may "ride along" and collect data at the same time if the data is useful to them. There are some operational restrictions to riding along; for instance, the two Microwave Remote Sensing instruments (RADAR and Radio Science) are both major power consumers and cannot be operated simultaneously.

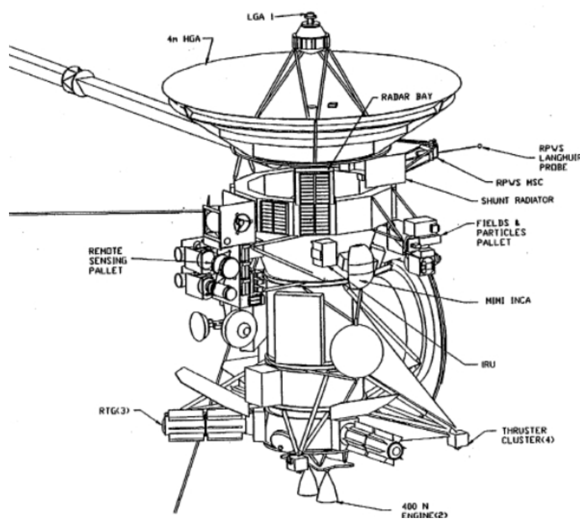

Figure 1. The Cassini Spacecraft.

In 2010 the Cassini Project completed tour planning for an additional 7-year phase called the Cassini Solstice Mission (CSM) that will extend the mission lifetime through Saturn's northern summer solstice. This extension permits observations of seasonal change across nearly half a Saturnian year (see Fig. 2).

II. How Cassini Plans Science

The process to plan CSM science started in January 2009 with the selection of a trajectory by the Cassini Project Science Group (PSG). This group, which meets three times a year, is made up of the Principal Investigators and Team Leaders of the 12 science instruments, interdisciplinary scientists, science planners, and various scientists from each instrument. Navigation engineers designed multiple trajectories that attempted to meet all the science objectives that the PSG selected for the CSM, targeting the planet, its rings, magnetosphere, icy moons, and Titan.

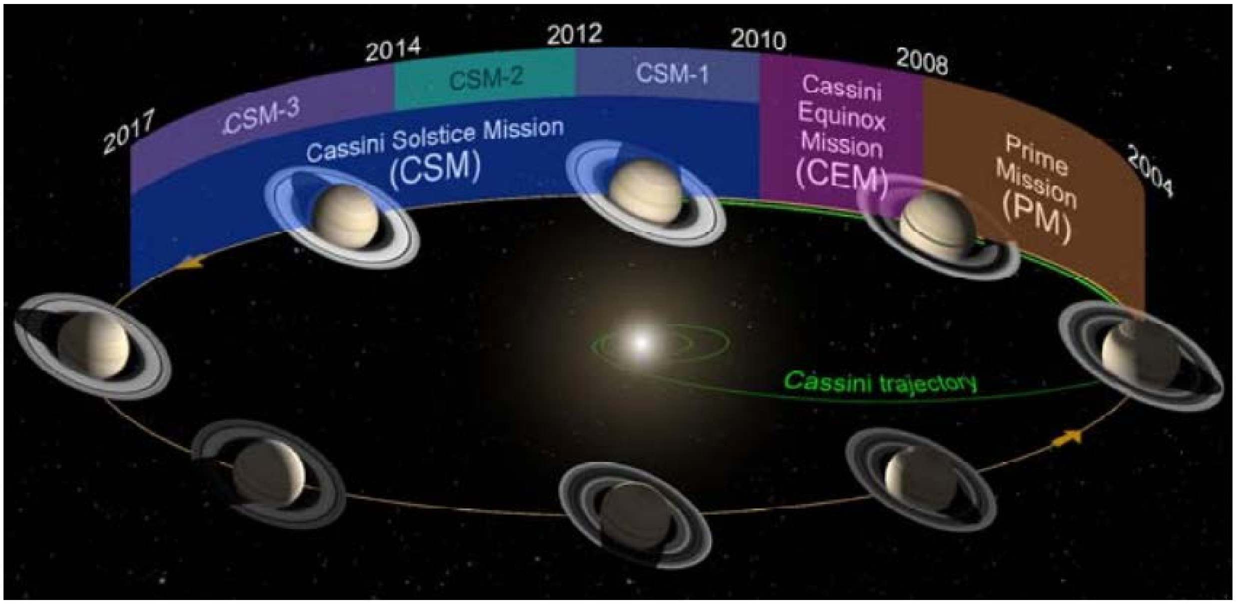

Figure 2. Phases of the Cassini Mission. The Cassini Prime Mission ran from just after winter solstice to just before Equinox. The Equinox Mission was centered on Equinox; the Solstice Mission will last until Saturn's summer solstice.

Once the trajectory was selected there was a very short period of three months during which the science community could request small changes ("tweaks") to the trajectory to improve science opportunities. The navigators accommodated these changes where possible. The PSG evaluated this revised trajectory, looking at how the proposed changes affected overall science opportunities and propellant use. Following this evaluation, the PSG decided which tweaks would become part of the final, official trajectory, named SM-7a[3].

The chosen trajectory contained a wealth of competing multi-disciplinary science opportunities (see Fig. 3). Making the most of these opportunities presented challenges in allocating observing time to different disciplines and instruments, and in preserving the precise timing required when there can be a gap of years from science selection to execution. Fairly allocating observing time among the disciplines required intense advance planning, complicated by needing consensus among the various disciplines. To accommodate all of these concerns, the science planning process was segmented along science discipline lines[4]. After the release of the final trajectory, Science Planning divided the entire trajectory into smaller segments that were assigned to science discipline working groups. There are six discipline working groups, made up of science planning engineers, scientists from instrument teams, and interdisciplinary scientists. Each working group focuses on a different aspect of Cassini science: the Titan Orbiter Science Team (TOST) concentrates on Titan observations, the Satellite Orbiter Science Team (SOST) on observations of all other satellites, and the Saturn and Rings Target Working Teams (TWTs) are responsible for Saturn and the ring system, respectively. The Magnetosphere and Plasma Science (MAPS) TWT focuses on Saturn's magnetosphere, while the Cross Discipline TWT considers all science objectives during apoapse periods. Each TWT or OST's segments include opportunities especially of interest to that TWT or OST. For example, TOST segments generally run from a day before each Titan encounter closest approach to a day after.

The science observations contained in each discipline segment must be considered against one more metric. CSM funding levels will be significantly lower than prime and extended mission funding. Consequently, all CSM science is driven by a carefully honed set of prioritized science objectives. To establish these objectives, each discipline working group identified their top priority science objectives for the CSM. These objectives either *i)* addressed the goal of observing seasonal change in the Saturnian system, understanding underlying processes, and preparing for future missions, or *ii*) were new questions that arose out of prime and extended mission science (e.g.

determining the composition and distribution of Titan's newly discovered lakes). The Project Scientist then constructed a matrix of CSM science objectives, and slotted them as Priority 1, 2, or 3. Each discipline working group must ensure that any observations planned for the CSM meet the Priority 1 objectives and as many Priority 2 and 3 objectives of the CSM as is possible within the observation time allotted to that discipline. Each working group is responsible for developing fully integrated timelines of the science that will be accomplished during their segments. Fully integrated segments are delivered to the Science Planning Team, which combines the segments into 10 week sequences that are uplinked to the spacecraft at the end of the implementation process. Each discipline developed its own method of getting from raw unintegrated segments to detailed designs ready to be included in a sequence. TOST's method is detailed in the next section.

.

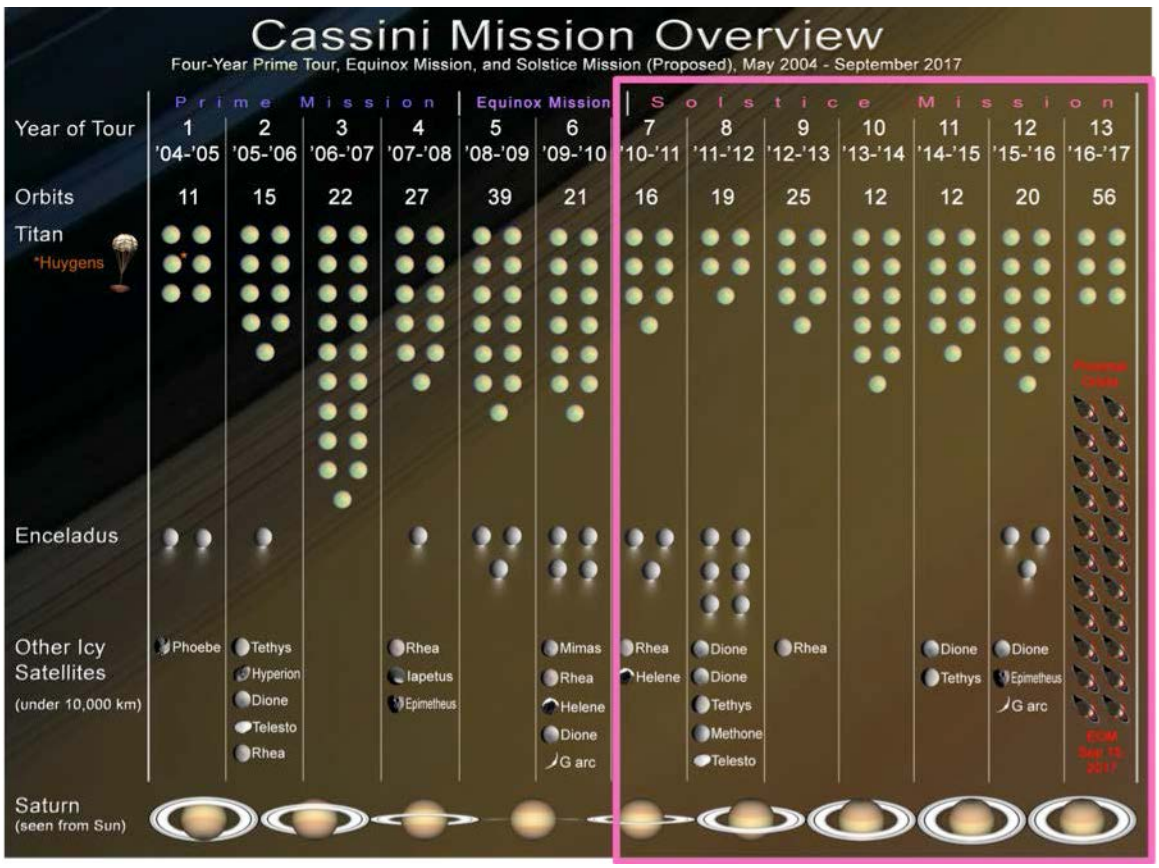

Figure 3. Cassini Mission Overview. A yearly summary of the number of orbits and flybys for the entire mission from 2004 through 2017. Figure courtesy David Seal, NASA/JPL

III. How TOST Plans Science

The process of integrating Titan flybys got an early start with a so-called "jumpstart", which maximized Titan science return by balancing science objectives across the entire set of flybys[4]. The major resource allocated during the jumpstart was pointing control of the spacecraft. The ORS instruments are mounted in fixed position to the spacecraft bus, so if one camera wants to point at Titan it is not possible to simultaneously point the HGA (for RADAR and RSS) at the same target. Consequently, been the designated "prime" instrument can be highly contested, and is the first issue to be worked out in any Cassini timeline. At the conclusion of the process, TOST had a package of master timelines (see e.g. Fig. 4) assigning observing and turn time for every TOST segment. The package was put on the shelf in this state until several months before the segment will be needed for sequence development.

The next step in the process leading to uplinkable flight sequences is *integration*. The high-level timeline for a flyby would be fleshed out (Figure 4). Power use ("op modes") and telemetry modes would be specified. The attitude strategy would be developed in greater detail, with testing to make sure that spacecraft turns to orientations near science targets of interest ("waypoints") and to downlink data could be executed safely and within allocated time. DSN station availability would be confirmed, and additional passes might be added to the timeline if needed for e.g. support of Radio Science Subsystem (RSS) science or a second playback of especially high-priority science. Instruments would request and negotiate data volume, ensuring that there were no storage or downlink capability issues.

Even with all of these details, the overall flyby timeline remained sacrosanct. With rare exceptions (such as when a trajectory change moved an occultation observation), integration of Titan flybys did not change the allocations of time when an instrument was in charge of determining the spacecraft pointing.

It is against this background that we now consider how TOST accommodated unexpected late changes in the integration phase to Titan science.

138TI_T72	8175	Lit Outbound				
Start Time	**End Time**	**Prime Activity**	**Obs. Detail**	**Op Mode**	**TLM Mode**	**Comments**
2010-267T05:17:00	2010-267T05:57:00	SP Turn to WP				
2010-267T05:57:00	C/A - 12:26:00	OD Uncertainty Dead Time				
C/A -12:26:00	-09:00	CIRS	N			
-09:00	-05:00	CIRS	R			
-05:00	-02:15	CIRS	T			
begin custom period						
-02:15	-00:15	CIRS	CIRS will turn to VIMS attitude			FIRLMB at 87S and 60N; FIRLMBAER and INT
-00:15	0	VIMS				
2010-267T18:38:41		CLOSEST APPROACH	NEG_Y to Titan,			
0	+02:15	VIMS				NT2 5km/pixel equat
end custom period						
+02:15	+05:00	UVIS	X			
+05:00	+09:00	UVIS	X			
+09:00	+14:00	VIMS	V ISS riding along			
+14:00	+22:28	VIMS	B ISS riding along			
C/A + 22:28:19	2010-268T17:22:00	OD Uncertainty Dead Time				
2010-268T:17:22:00	2010-268T18:02:00	SP Turn to Earth for downlink				
2010-268T18:02:00	2010-268T19:32:00	Y-Bias window				
2010-268T19:32:00	2010-269T06:17:00	Canberra 70M Array		DFPW Normal		

138TI_T72	8175	Lit Outbound				
Start Time	**End Time**	**Prime Activity**	**Obs. Detail**	**Op Mode**	**TLM Mode**	**Comments**
2010-267T05:17:00	2010-267T05:57:00	SP Turn to WP	NEG_Y to Titan, NEG_X to NTP	DFPW Normal	S_N_ER_3	
2010-267T05:57:00	C/A- 12:26:00	OD Uncertainty Dead Time		DFPW Normal	S_N_ER_3	
C/A -12:26:00	-09:00	CIRS	N	DFPW Normal	S_N_ER_3	
-09:00	-05:00	CIRS	R	DFPW Normal	S_N_ER_3	
-05:00	-02:15	CIRS	T	DFPW Normal	S_N_ER_3	
begin custom period				DFPW Normal	S_N_ER_3	
-02:15	-00:15	CIRS	CIRS will turn to VIMS attitude	DFPW Normal	S_N_ER_3	FIRLMB at 87S and 60N; FIRLMBAER and INT
-00:15	0	VIMS		DFPW Normal	S_N_ER_3	
2010-267T18:38:41		CLOSEST APPROACH	NEG_Y to Titan, (Tc2a)	DFPW Normal	S_N_ER_3	
0	+02:15	VIMS		DFPW Normal	S_N_ER_3	NT2 5km/pixel equat
end custom period				DFPW Normal	S_N_ER_3	
+02:15	+05:00	UVIS	X	DFPW Normal	S_N_ER_3	
+05:00	+09:00	UVIS	X	DFPW Normal	S_N_ER_3	
+09:00	+14:00	VIMS	V ISS riding along	DFPW Normal	S_N_ER_3	
+14:00	+22:25	VIMS	B ISS riding along	DFPW Normal	S_N_ER_3	
C/A + 22:25:00	2010-268T17:22:00	OD Uncertainty Dead Time		DFPW Normal	S_N_ER_3	
2010-268T17:22:00	2010-268T18:02:00	SP Turn to Earth for downlink		DFPW Normal	S_N_ER_3	
2010-268T18:02:00	2010-268T19:32:00	90 min Y-bias window		DFPW Normal	S_N_ER_3	
2010-268T19:32:00	2010-269T06:17:00	Canberra 70M Array		DFPW Normal	RTE_N_SPB	

Figure 4. Example of master timeline for a Titan flyby at the end of the jumpstart process (top) and after full integration (bottom). This master timeline for T72 gives a time ordered listing of which team controls the spacecraft attitude at every point in the flyby period. Times are given in absolute spacecraft time or in flyby closest-approach epoch-relative time. Templates are noted under observation details. Operational modes and telemetry modes are not completed until detailed integration immediately prior to sequence development.

IV. Incorporating Changes Example A: Stacking new science on top of old (RSS LGA)

Like many missions, Cassini has long used the concept of "rider" observations that piggyback on another instrument's science. Generally, this meant that one ORS instrument optimized pointing for their FOV and the remaining ORS instruments would collect data using the prime pointing. The RSS was unable to take advantage of ride-along opportunities due to its highly specific pointing requirements utilizing the HGA. As planning commenced for the last few years of the mission, the RSS scientists wondered: would it be possible to use the spacecraft's LGA, and use its more tolerant pointing requirements to allow RSS to ride along with other Cassini science?

A. Gravity Science Needs

RSS Titan gravity flybys are used to determine if the moon has an internal ocean. A typical gravity flyby may last from 24 to 32 hours, during which time the spacecraft's HGA is turned towards Earth. Multiple gravity flybys are needed to accurately determine the moon's geoid and Love number. Initially, RSS was allocated five Titan gravity flybys during the Prime and Extended mission. Determining Titan's precise geoid proved to be more difficult than initially expected, so RSS requested five additional gravity flybys in the Solstice mission. TOST awarded RSS three gravity flybys based on simulation results showing that three flybys would be sufficient to meet the RSS science goals.

The RSS team argued that additional flybys were needed, and began examining the feasibility of using a low-gain antenna to gather additional gravity data without the need for a dedicated flyby. Although the LGA tracks would not be of the same high caliber as the HGA tracks, the LGA tracks would provide adequate data, improve latitude-longitude coverage, and therefore increase the likelihood of determining the existence of a Titanian subsurface ocean.

Table 1. RSS LGA Opportunities: Implementation Status

	Goal	Target	Floor	Comments
Altitude	1800 Km	1600 to 2500 km	<4000 km	Balance desire for close flyby for determining higher degree gravity field components with need to be out of Titan's atmosphere
SEP Angle	Larger is better	>45 degrees	>30 degrees	Mitigate effects of solar plasma
Lat/Long	Variety	Variety	Variety	Crucial for geoid determination
S/C–Earth visibility	Continual clear LOS for +/- 4 hours and on wings			+/- 2 hours, some holes in coverage acceptable
DSN elevation angle			>25 degrees	As large as possible

The first need was for the RSS scientists to develop requirements for which flybys could be utilized for LGA gravity science. Criteria (see Table 1) included the flyby altitude, Sun-Earth-Probe (SEP) angle, a sufficiently high DSN station angle at closest approach, no Earth occultation at closest approach, and only using flybys on reaction wheels (flybys on thrusters would be too dynamically noisy for gravity science). RSS presented their work to Cassini's Titan Working Group, which approved the proposal.

The team also developed DSN tracking requirements for supporting the LGA science. Ideally, there would be 2-way DSN tracking centered within +/- four hours of closest approach, with a minimum requirement of +/- two hours. Though 70 meter coverage would be needed for the closest approach interval, the wings away from closest approach could utilize 34 meter tracking.

The next step was testing the concept on the spacecraft to determine if Cassini's LGA-1 antenna (co-aligned with the HGA) could gather adequately precise Doppler tracking. The test, conducted in the summer of 2010, confirmed for the RSS team that scientifically-useful Doppler data could be obtained with the LGA moving over a 55 degree cone angle off of Earth point.

Next, RSS designed a ranking system to assess the suitability of each remaining Titan flyby for RSS LGA science using the criteria listed above. The team ended up with a list of nine flybys that would be good candidates for successful LGA science.

B. Integrating RSS LGA Science Opportunities

Once the science and feasibility issues were addressed, the science planners turned to addressing how to integrate and implement the RSS LGA science opportunities.

During integration, science planners needed to submit initial requests for spacecraft and ground resources. This included DSN station requests needed to support the RSS LGA opportunity (70m stations near the closest approach period and 34m requests on the wings) and the RSS science request itself. The RSS would not specify spacecraft attitude, but would request that the instrument controlling spacecraft pointing use the NEG_Z axis (along which the LGA-1 was oriented; see Figure 1) as the secondary axis. By negotiation with the rest of the Titan science team, this was *not* considered a mandatory constraint on the spacecraft pointing—if the prime instrument could not meet its own science needs with the RSS-friendly orientation, the prime instrument was free to choose a different orientation. However, detailed pointing design wouldn't be performed until the start of the sequence implementation phase. This meant that the RSS team would not know until implementation work started if the pointing design would actually be able to support the RSS LGA science.

C. Implementing RSS LGA Science Opportunities

One of the first actions in Cassini's sequence implementation is the early delivery of the final pointing designs for the prime observations, a week before the first ("port 1") merge. These early deliveries are only required if the prime observation will be used by any collaborative rider instruments, so that the prime and rider teams have an opportunity to work together to create a design that works for both teams. Though as already mentioned the RSS observation cannot compel the prime pointing observation to redesign to accommodate RSS, this early delivery window is used so that RSS can perform analysis to determine if the flyby meets the criteria for a successful LGA gravity flyby. If the RSS analysis shows that the pointing cannot support RSS LGA science, then the RSS LGA science request is withdrawn. Table 2 indicates which flybys to date have not continued on in implementation. DSN passes in support of LGA science are released at this point so they can be used by other projects.

If RSS analysis shows that the LGA science is supportable, the project schedules a preliminary approval meeting between the Sequence Integration Process (SIP) leads, project management, and the RSS, SCO, and NAV teams. This is the point where SCO and NAV would see if the proposed LGA science meets the engineering and navigation "go" criteria. As Table 2 shows, only one proposed RSS LGA science opportunity—on Titan flyby T97--has reached this preliminary approval stage. Though the SCO and NAV teams gave their approval, project management decided that the LGA science was merely going to replicate the RSS HGA science from the upcoming T99 flyby, and thus implementing the RSS LGA science on T97 was not worth the risk. The SIP leads removed the RSS science from the sequence, and released the associated DSN passes.

The final step in RSS LGA implementation would be a final approval meeting where all teams would provide "go" authority. By this point, SCO would build and test a real-time file containing the LGA commands.

Table 2. RSS LGA Opportunities: Implementation Status

Flyby	C/A Date	Sequence	OK in Integration?	Pointing Design Compatible?	Project Review Decision
T82	Feb 19 2012	S72	Yes	No	n/a
T94	Sept 12 2013	S80	Yes	No	n/a
T97	Jan 1 2014	S82	Yes	Yes	No go
T105	Sept. 22 2014	S85	Yes	TBD	—
T110	Mar 16 2015	S88	TBD		
T111	May 7 2015	S89	TBD		
T115	Jan 16 2016	S92	TBD		
T116	Feb 1 2016	S92	TBD		
T123	Sept 27 2016	S96	TBD		

D. Stacking new science on top of old: bottom line

As this example shows, it is indeed possible to create additional piggybacked science opportunities for an in-flight mission. However, it may require a significant amount of work, time, and coordination. As Table 2 shows, there are four more possible RSS LGA opportunities remaining in the Cassini mission. We are hopeful that we will be able to report that at least one has been successfully implemented, showing that RSS gravity science can indeed use the LGA option with non-optimal pointing.

V. Incorporating Changes Example B: Accommodating a damaged instrument

Losing science instruments is a risk for any mission, especially those operating long past prime mission. If an ailing instrument may or may not be restored to operation, how should science planning and operations respond? This section addresses one approach to pushing timeline integration decisions as late as possible.

A. CAPS anomaly

The Titan jumpstart process allowed each instrument team to designate the two flybys which would offer the most unique/important science for the instrument. A "10-pointer flyby" quickly became project shorthand for high-priority science. AACS analysts would avoid proposing changes in spacecraft orientation during these observations. If another project demanded DSN coverage that was needed to downlink this critical high-priority data, saying "10 pointer" conveyed the urgency and uniqueness of the data.

So it was especially bad news to learn in mid-June of 2011—six months before a CAPS 10-pointer Titan flyby--that a spacecraft voltage anomaly was caused by the CAPS instrument and that the instrument was being turned off pending a tiger team investigation. CAPS and TOST were hopeful that the anomaly could be resolved in time to have CAPS back online for the T79 flyby in December 2011. However, the S71 sequence, which contained the T79 flyby segment, was already in implementation. Final approval of the sequence was scheduled for November 8, less than four months away. If the project gambled that CAPS would be operational by December, they risked having a Titan flyby which gathered little science if CAPS remained off. If, on the other hand, the flyby time was reallocated to another instrument under the assumption that CAPS wouldn't recover in time, we risked losing a unique CAPS science opportunity if the instrument did turn back on before the flyby.

By late July, the project and the CAPS team had decided that CAPS would not be turned back on in the immediate future; long-range operations plans for the instrument were still pending a tiger team decision. Together, the project, TOST and the sequence implementation leads decided to develop two possible timelines for the sequence. One would include CAPS science as originally planned. The other alternate timeline would replace the CAPS science with observations from another instrument. This alternate high-level plan needed to be quickly developed by TOST—within a matter of days—in order to support the analysis and testing that is a normal part of the sequence development process.

B. Creating and managing alternate timelines

The decision to create two alternate timelines created major logistical issues for the flight and instrument teams. Cassini's planning database was not created to manage developing and managing dual timelines during integration or implementation. CIMS, the database that Cassini uses to manage instrument observation requests, is not designed to handle more than one timeline for a given time period. All requests starting or ending within the boundaries of a segment or sequence would be assumed to belong to that segment or sequence. So, the TOST and sequence leads developed an alternative naming scheme that would be used for the alternative timeline. After verifying that the plan would work in the CIMS testbed development space, the leads then created a new sequence that contained all events from both the CAPS and alternative timeline, and then carefully hand-edited the new delivery to remove the CAPS observations. Like weeding a garden, this process was ongoing: every time that an observation from the T79 flyby was edited, the revised observation would show up in both the CAPS and alternative timelines. The sequence leads needed to ensure that observations were properly retained in the appropriate deliveries.

The next major issue was deciding if the flyby epoch would be the usual ground-moveable block (GMB, which can only be revised before the background sequence is radiated to the spacecraft) or a live moveable block (LMB, which can be revised until just prior to execution). Developing the flyby as an LMB would give more flexibility, allowing the Cassini team to delay making a decision between the two timelines up until a few weeks prior to the actual encounter. But using an LMB would add complexity: we would need two different LMB epoch times in order to more easily verify and validate separate timelines, the RBOT biases would need to be sent as real time commands rather than set within the background sequence since the bias strategy would be different for the two alternate timelines since their pointing strategies would be different. Most importantly, the sequence leads and instrument teams would need to check both timelines—delivering multiple files, performing multiple checks—for three input ports, adding significantly to workload. Choosing to use a GMB would force the project to decide on a specific timeline far sooner, in which case the team would only need to verify and validate two separate background sequences through one port. In the end, the project took the simpler less risky route and chose to use a GMB epoch, which was deemed less risky operationally because of the reduced timeframe for responding to change. There would be more opportunities to verify the sequence commanding using a GMB since decision was locked in with time for another two ports of analysis.

Asking the teams to support two alternate timelines through even one port was an undertaking. Each affected team needed to provide two separate input files, accounting for different activities, alternate naming conventions, and two different data volume strategies. Eight teams in all were impacted, including six of the science instruments (CIRS VIMS ISS UVIS CAPS RPWS), the AACS team, and Science Planning (which manages data volume). In addition, the pointing designs for the alternate timeline needed to be developed and verified on a highly compressed schedule. CIRS and VIMS (the two ORS instrument teams that would be prime on the alternate timeline) needed to do full PDT designs of their observations and test to make sure they are valid in less than two weeks, a task that normally would be completed within two to three months. These designs then had to be tested by the AACS team for RBOT compliance, squeezing in this work to confirm that both alternate timelines were workable in time for the so-called "Port 3" input port on September 12.

Our schedule called for the project to make a final decision by the end of September, a mere two months after the dual timeline development began. After careful study, the Cassini project scientist decided that at the time by which a decision had to be made it had not been determined if CAPS could be safely turned on without risking the spacecraft's health and safety. Nor was it likely that those studies would be completed in time for the T79 flyby in December 2011. The project decided to proceed with the alternative ORS timeline. While it was disappointing to lose the highly anticipated CAPS science, the alternative timeline work enabled the return of ORS Titan observations and maximized science return for the flyby.

C. Accomodating a damaged instrument: bottom line

The dual timeline development option represented a serious increase in workload for many members of the flight team. Creating a planning database that could easily manage alternate timelines would have made the process easier, albeit at a greater development cost. Missions that anticipate a greater likelihood of loss of capability leading to a need to replan science should consider this option.

VI. Incorporating Changes Example C: Dwindling mission consumables

Mission consumables are one of the main limitations on science planning. As a mission moves into late or extended status, using consumables may become easier (because e.g. usage has been much lower than expected, or because mission success objectives have already been fulfilled). In other cases, the mission management may decide to keep much closer rein on use of some consumables. This section addresses how the TOST group adapted when early flyby planning assumed more liberal guidelines than were present when the flybys were closer to implementation.

A. Tracking Hydrazine Use

As Cassini moved into its second extended mission phase (lengthening the mission by seven years), the project focused more intently on the use of mission consumables. Hydrazine, in particular, was an issue. Mission planning estimates showed minimal margin for completion of the proposed trajectory, so the project carefully examined all proposed hydrazine use to see where savings might be eked out. Science use of hydrazine—particularly that used to enable faster turn rates--came under scrutiny.

In particular, the project needed to look at the most accurate and detailed estimates of hydrazine use during each flyby. Cassini's AACS team produces these models, using the detailed pointing commands created by the instrument team responsible for spacecraft pointing while the spacecraft is on thruster control. However, during the jumpstart the instrument teams hadn't yet created the detailed pointing for their observations, so there were no high-fidelity estimates of hydrazine consumption. TOST provided estimates of hydrazine use for those flybys based on historical patterns. These estimates were then used to construct the hydrazine budget for the remainder of the mission.

Project management was concerned that they had limited options for dealing with higher-than-expected hydrazine use. Previously, the project only got the AACS hydrazine estimates prior to the so-called "port 3" phase of implementation. If hydrazine use was too high, it would be too late to redesign the observation; the project manager's only choice would be to pull the observation. This presented a significant problem for the instrument teams because it was too late to make smart science choices. Requesting the instrument teams to redesign their observation to use less hydrazine late in implementation would mean no time to redesign, leading to loss of an observation during the valuable closest approach interval. The project manager suggested that it would be easier to look at hydrazine usage earlier in the integration process.

B. Designing an early look system

The TOST leads decided to come up with a process to get an early look at hydrazine use. We decided to that we ask the instrument teams to deliver their detailed pointing designs in the integration phase, months earlier than normal. The AACS team agreed to run their hydrazine models on these early deliveries. A TOST science planning engineer would graph the AACS results, using the proposed pointing timeline to indicate the cost of each pointing command carried out in the proposed design (see Figure 5 below). If the proposed usage was too high, there was time to either modify the design so it used less hydrazine, or to give the observation time to another instrument or to trade with a later flyby. TOST scientists—who had integrated the flyby timeline—would make a science-based decision, not the Project Manager.

C. Dwindling Mission Consumables: bottom line

For the cost of a little more up-front work, instrument teams were able to have the best shot at ensuring their observations wouldn't be cut due to heavy consumables use. TOST would have the best shot at making sure that observing opportunities weren't lost. And the project would have the assurance that consumables were being managed effectively.

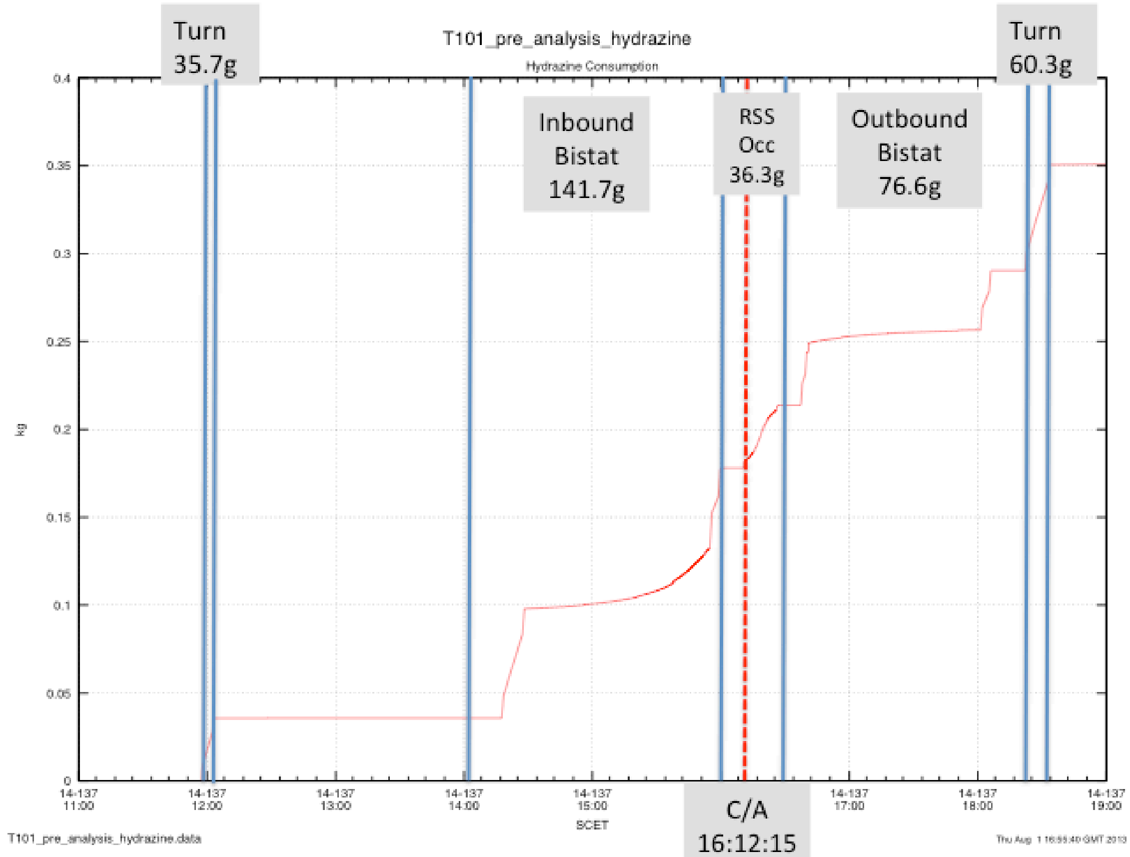

Figure 5. Example of hydrazine use breakdown by individual activity for a Titan flyby on thrusters. This timeline for the T101 flyby shows a cumulative trace of total hydrazine use as a function of time (red solid line). Individual activities—turns, RSS bistatic observations and occultations—are labeled with the amount of hydrazine consumed for each activity, with solid vertical blue lines showing their extent in time. The time of closest approach is shown as a vertical red dashed line.

VII. Incorporating Changes: One instrument's high priority science is another instrument's heating

As discussed earlier, one of the advantages of the jumpstart process was being able to allocate Titan science observations across the entire Solstice mission, so all science trades would be completed up front. Even so, detailed observation designs uncovered some surprises where one instrument's plans would seriously impact another instrument's opportunities. This section examines how the Titan scientists and integration team were able to address those challenges.

A. RADAR high priority observation heats VIMS

During the jumpstart process, RADAR claimed the T104 flyby as one of its top priority "10 pointer" opportunities. The groundtrack and timing made the flyby an excellent candidate for detecting seasonal changes to seas and empty lakes.

Unfortunately, early analysis of RADAR's planned pointing showed that the proposed observation would heat the VIMS instrument by 6 degrees K, placing the heating event at consumable level. Heating VIMS or CIRS at non-consumable levels is common, and by this "extended-extended" mission phase even consumable-level heating is acceptable. Generally, there is enough time between the heating event and any subsequent VIMS or CIRS observations for the instruments to cool back down to normal operating temperature. Unfortunately, for T107 the RADAR-caused heating would obliterate a VIMS observation scheduled for immediately after RADAR's observation. The VIMS team had also discovered on two previous flybys that their heating model was underpredicting actual heating levels by 2-3 degrees K, making the team less likely to approve predicted heating. VIMS scientists felt strongly that their proposed science was important, and asked RADAR to redesign their observation to avoid high levels of VIMS heating.

In response, the RADAR team developed three additional observational strategies. The initial plan was removed from consideration, especially as one possible implementation of the RADAR observation would send the predicted VIMS heating to near the absolute flight-rule limit of +15.75 degrees K. RADAR compiled a presentation describing each option, including the science enabled by each strategy, and the resultant VIMS and/or CIRS heating. All instrument teams were invited to comment on how the proposed alternative designs would impact their observations.

It became clear that the discussion would be more involved than could be accommodated at the regular integration meeting, so we scheduled a special core science meeting, with one representative from each science team. Each team that was affected—VIMS, RADAR, and INMS, which was riding along on the RADAR observations—presented what science they wanted, and the impact on that science of each proposed option. As a group, a final design was chosen that did include as expected significant VIMS and CIRS heating, but which provided some limited ability for VIMS to collect data even after heating. VIMS pointed out that there was a later flyby that would help them get what they wanted T108 VIMS followed by RADAR

B. High priority science/heating tradeoff: bottom line

The successful resolution of the conflict illustrated the virtues of having instrument teams work together to find the best overall science balance. Our scientists really know a lot about each others' instruments; for example, ISS was able to point out VIMS detection of specular reflection—a high priority for VIMS-- could still be implemented for the T104 flyby.

VIII. Conclusion

12 science instrument teams, including agreement on what science would be accomplished during each flyby. By looking at all 56 flybys at once, the best balance of interior, surface, atmospheric, and magnetospheric interaction science was achieved. By deciding on the closest approach attitudes early, it was possible to influence the final trajectory production and change some flyby altitudes to improve scientific return. In less than 15 hours of teleconference time, integrated conflict-free timelines were completed for each Titan flyby detailing allocation of the time outside closest approach using re-useable templates. By completing the jumpstart during the equinox mission which is funded at the same level as the prime mission, the TOST team was able to take advantage of full participation by key long-range-planning personnel who might not be able to attend as many meetings during the CSM due to the lower funding profile. This process allowed the Cassini mission to maximize Titan science return across the CSM.

Though early and comprehensive high-level planning of science timelines has its benefits, the resulting plans may not offer flexibility for late changes. Though it is difficult to know in advance specifically *which* changes will be needed, certain *types* of changes can be expected. The examples described out in this paper—depleted mission consumables, balky science instruments, conflicts between two desired instrument operational environments and a desire to add more activities to an existing timeline—are typical. What is unique is being able to develop ways to accommodate these changes.

Acknowledgments

Thanks to William Heventhal III, Jan Berkeley, Anna Marie Aguinaldo for useful discussions. This effort was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration

References

[1]Buffington, B., Strange, N., and Smith, J., "Overview of the Cassini Extended Mission Trajectory," 26th AIAA/AAS Astrodynamics Specialist Conference, paper AIAA-2008-6752, August 18-21, 2008.

[2]Paczkowski, B.G., Larsen, B.S., and Ray, T.L., "Managing Complexity to Maximize Science Return: Science Planning lessons Learned from Cassini", IEEE Aerospace Conference, Big Sky, Montana, March 7-14, 2009.

[3] Smith, J., and Buffington, B., "Overview of the Cassini Solstice Mission Trajectory," AIAA/AAS Astrodynamics Specialist Conference, Pittsburgh, PA, AAS Paper 09-351, August 10-13, 2010.

[4]Steadman, Kimberly B., Pitesky, Jo E.., Ray, Trina L., Burton, Marcia E., and Alonge, Nora K., "Cassini Titan Science Integration: Getting a "Jumpstart" on the Process", AIAA SpaceOps Conference Proceedings, April 25-30, 2010.

[5]Paczkowski, B.G., "Cassini Science Planning Process", AIAA SpaceOps Conference Proceedings, May 17–21, 2004.

[6]Hansen, C.J., Waite, J.H., and Bolton, S.J., "Titan in the Cassini-Huygens Extended Mission", *Titan from Cassini-Huygens* edited by Robert H. Brown, Jean-Pierre Lebreton and J. Hunter Waite, Springer Science+Business Media, London, 2009, pp. 455–478.