

# INITIAL CASSINI PROPULSION SYSTEM IN-FLIGHT CHARACTERIZATION

T. J. Barber and R. T. Cowley\*  
Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, CA

## Abstract

The Cassini Propulsion Module Subsystem (PMS) has performed excellently throughout the first four years of mission operations. The PMS is the most complex interplanetary propulsion subsystem ever flown, with separate monopropellant and bipropellant propulsion modules, each replete with many redundant components. Propulsion system performance has generally been exemplary, with the apparent exception of large leakage through the primary hard-seat regulator. Maneuver performance has been excellent, and the spacecraft has flown a complex, Venus-Venus-Earth-Jupiter gravity assist trajectory virtually trouble free. There is no early evidence for pressure transducer drift, in contrast with the Galileo, Voyager, and TOPEX-Poseidon missions. PMS system health checks based on helium mass calculations (helium mass budgets) have demonstrated no discernible helium pressurant or propellant leakage to date. An extensive, in-flight assessment of 0.9-N monopropellant hydrazine thruster performance during pulse-mode operation has been undertaken. Main engine and hydrazine thruster performance has been within specification. Propellant and consumable usage have generally been lower than predicted, and the prospect for a highly successful, multi-year scientific investigation of the Saturnian system remains excellent.

## I. Introduction

The Cassini mission to Saturn was conceived over a decade ago as a follow-up to the highly successful Pioneer 11 and Voyager missions to the ringed planet. Building upon the legacy of the Galileo mission to Jupiter, Cassini will enable an extensive, four-year investigation of Saturn's atmosphere, rings, icy satellites, magnetic field, and intriguing moon Titan. The Cassini mission is a joint NASA/ESA endeavor, and the Cassini orbiter is transporting ESA's Huygens probe, which will explore the atmosphere and possibly

the surface of Titan *in situ*. The Cassini mission is managed by Caltech's Jet Propulsion Laboratory (JPL) for NASA's Office of Space Science.

Cassini was launched on October 15, 1997, from Cape Canaveral, Florida, on a Titan IV-B launch vehicle. A Centaur upper stage placed Cassini on the proper interplanetary trajectory, a Venus-Venus-Earth-Jupiter Gravity Assist (VVEJGA) path to Saturn (see Figure 1). The selected trajectory is very similar to Galileo's Venus-Earth-Earth Gravity Assist (VEEGA) trajectory.<sup>1</sup> Cassini flew by Venus for the first time on April 26, 1998, obtaining the first of four required gravity assists in the VVEJGA trajectory. A large (450 m/s) plane-change maneuver known as the Deep Space Maneuver (DSM) was executed successfully on December 3, 1998. This set up the proper arrival conditions for the second Venus flyby on June 24, 1999. An Earth gravity assist flyby followed a mere seven weeks later, on August 18, 1999. The spacecraft traversed the asteroid belt and performed its final gravity assist at Jupiter on December 30, 2000. The Jupiter flyby offered an excellent chance for science instrument checkout and calibration. In fact, the simultaneous Cassini/Galileo observations of the Jovian magnetic field offered a unique science opportunity.

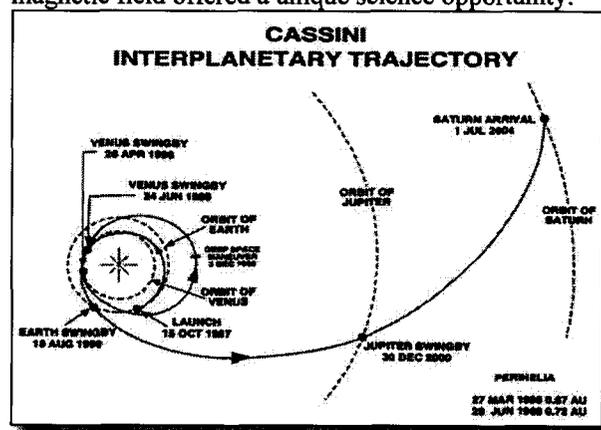


Figure 1. Cassini Heliocentric Trajectory

\*Members, Technical Staff

Cassini remains in excellent health in the outer solar system. Following a close flyby of Saturn's enigmatic moon, Phoebe, in mid-June of 2004, Cassini will arrive at the ringed planet on July 1, 2004. The spacecraft will fire a main engine to become the first artificial satellite of the planet Saturn. This Saturn Orbit Insertion (SOI) maneuver will be, coincidentally, a carbon copy of DSM—a ninety-minute, main-engine, pressure-regulated burn. A diagram of the Saturn arrival geometry and initial few orbits is presented in Figure 2.

Following SOI, the spacecraft will perform another large, main-engine maneuver to raise the orbit periapsis. This Periapsis Raise Maneuver (PRM) will occur on September 7, 2004. This maneuver places the entire spacecraft on a Titan-impact trajectory to prepare for the late-December 2004 Huygens probe release. Huygens will fly ballistically into Titan's atmosphere, while the Cassini orbiter will be deflected with a final pressure-regulated, main-engine burn called the Orbiter Deflection Maneuver (ODM). In mid-January of 2005, the Cassini orbiter will fly above Huygens as it descends through Titan's atmosphere over the course of a few hours, perhaps surviving landing on a solid or even liquid hydrocarbon surface (see Figure 3 for the Huygens probe entry and descent profile). Titan is enshrouded in clouds and high-altitude hazes; though its average surface temperature is a frigid 95 K, it is of high scientific interest, given the presence of organic molecules that might resemble pre-biotic conditions on Earth.

Following the completion of the Huygens phase of the mission, Cassini will execute a complex orbital tour of Saturn over four years (see Figure 4). More than seventy orbits of Saturn will be executed, the majority of which will include close Titan flybys. Titan is essentially the only Saturnian moon large enough to be used for gravity assist, which is required to modify the orbital trajectory. This technique was used with great success during the Galileo orbital tour of Jupiter.<sup>2</sup> A radar-mapping instrument on the Cassini orbiter will penetrate the clouds of Titan. This will allow the first-ever surface map to be generated of this intriguing world, similar to Magellan's unveiling of the planet Venus in the early 1990's. Cassini will likely have resources for an extended mission, allowing further characterization of this miniature solar system nearly a billion miles from Earth.

## II. The Spacecraft (Orbiter)

Cassini is the most complex interplanetary spacecraft ever built, likely representing the end of an era for NASA. The orbiter is a three-axis stabilized spacecraft with a launch (stack) mass of nearly six metric tons. Redundant components are an important part of the Cassini spacecraft design, in all subsystems. Table 1 contains some key "metrics" that characterize the Cassini spacecraft.

Figure 5 shows the spacecraft cruise configuration after launch vehicle adapter release and before probe release. The long Cassini magnetometer

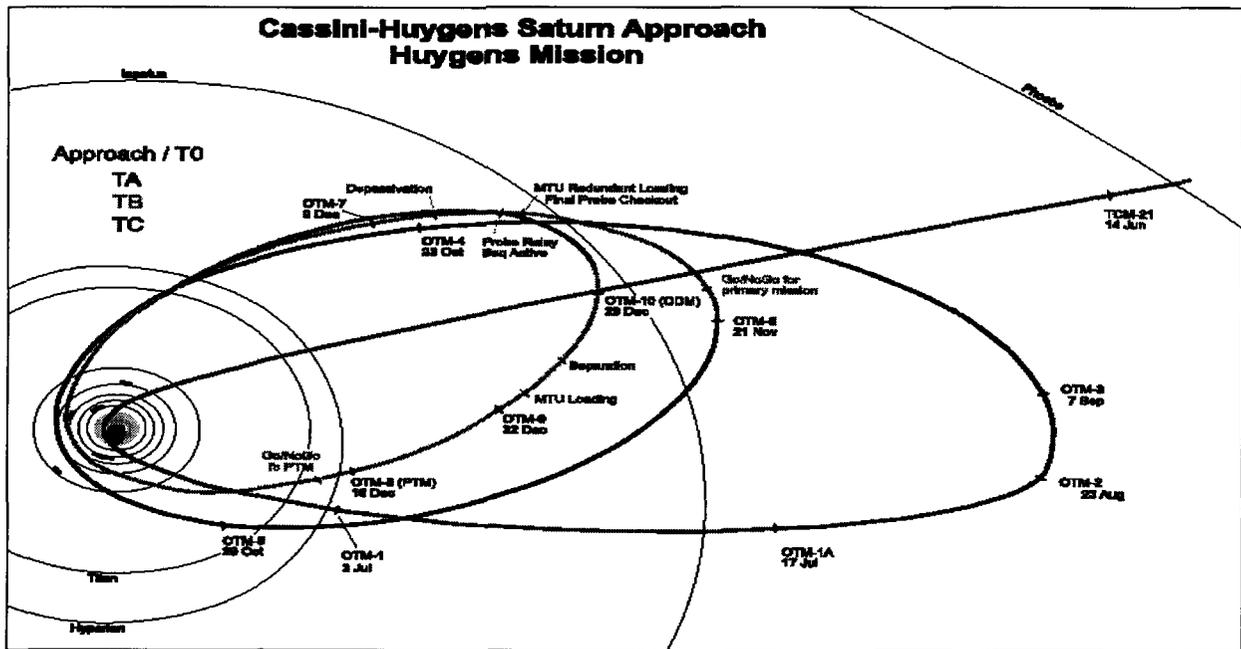


Figure 2. Cassini Approach and Early Tour Geometry

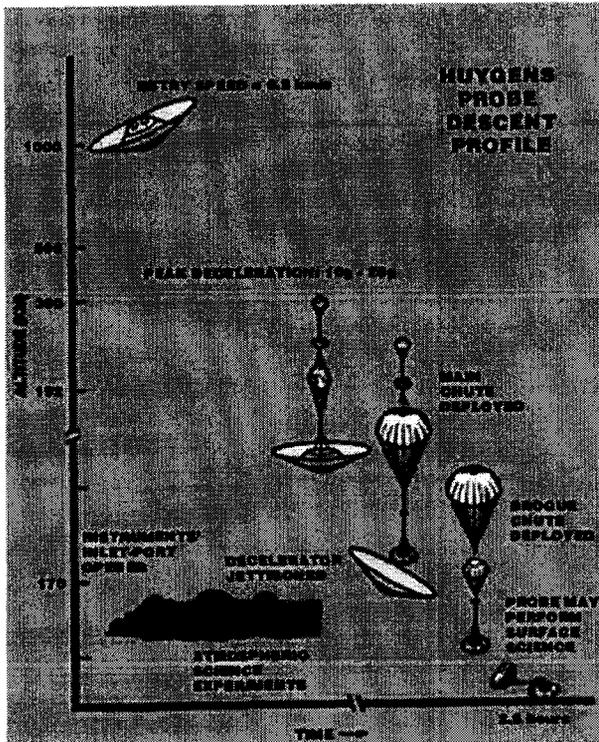


Figure 3. Huygens Descent Profile

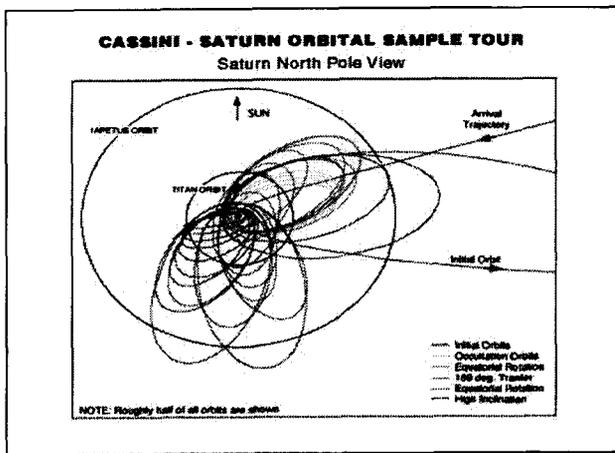


Figure 4. Cassini Orbital Tour Petal Plot

boom is omitted for clarity; in flight, the magnetometer boom was deployed without incident just after the Earth flyby in August of 1999. Similar deployment events occurred nominally just after launch, primarily science instrument cover deployments.

Due to the weak solar intensity at Saturn ( $< 16 \text{ W/m}^2$  on average), the orbiter is powered using three Radioisotope Thermoelectric Generators (RTGs). The

Table 1. Cassini Spacecraft Physical Characteristics

Parameter	Value
In-Flight Height	6.8 m
In-Flight Span (excluding Mag. Boom)	4.0 m
Centerline to Magnetometer Boom Tip	11 m
Spacecraft Initial Mass	5574 kg
Orbiter Dry Mass (no propellant/He)	2113 kg
Huygens Probe Mass	320 kg
Total Bipropellant Mass	3000 kg
[Oxidizer Mass]	[1869 kg]
[Fuel Mass]	[1131 kg]
Bipropellant Helium Mass	8.6 kg
Total Monopropellant Mass	132 kg
Monopropellant Helium Mass	0.4 kg
Launch RTG Power	879 W
End-of-Mission RTG Power	687 W
<u>Computer Resources</u>	
1750-A Flight Computers	6
Memory	512 KB
Solid State Recorder Storage	256 MB
<u>Telecommunication Resources</u>	
Maximum Engineering Data Rate	1896 bps
Maximum Science Data Rate	166 kbps

total RTG power output decreases from 879 W at the beginning of the mission to 687 W on the projected prime mission completion date. The power subsystem has worked quite well in flight; RTG power predictions have typically been within a few watts of actual values.

The Command and Data Subsystem (CDS) of Cassini represents a significant improvement vs. Galileo. In particular, the computer processing power was increased greatly, and the data storage medium was updated to a solid-state recorder. Galileo used a reel-to-reel tape recorder, which experienced some anomalous behavior in flight shortly before Jupiter arrival. Of note, much of Cassini's flight software work was deferred until after launch, given the long trip time to Saturn and budget pressures prior to launch. Cassini's CDS is fully redundant, and has worked very well in flight.

Cassini's telecom subsystem utilizes a fixed high-gain antenna (HGA) and two low-gain antennas (LGAs). Two LGAs were required for redundancy and for communication with Earth during the early portion of the mission. Specifically, spacecraft thermal constraints required the HGA to be sun-pointed during

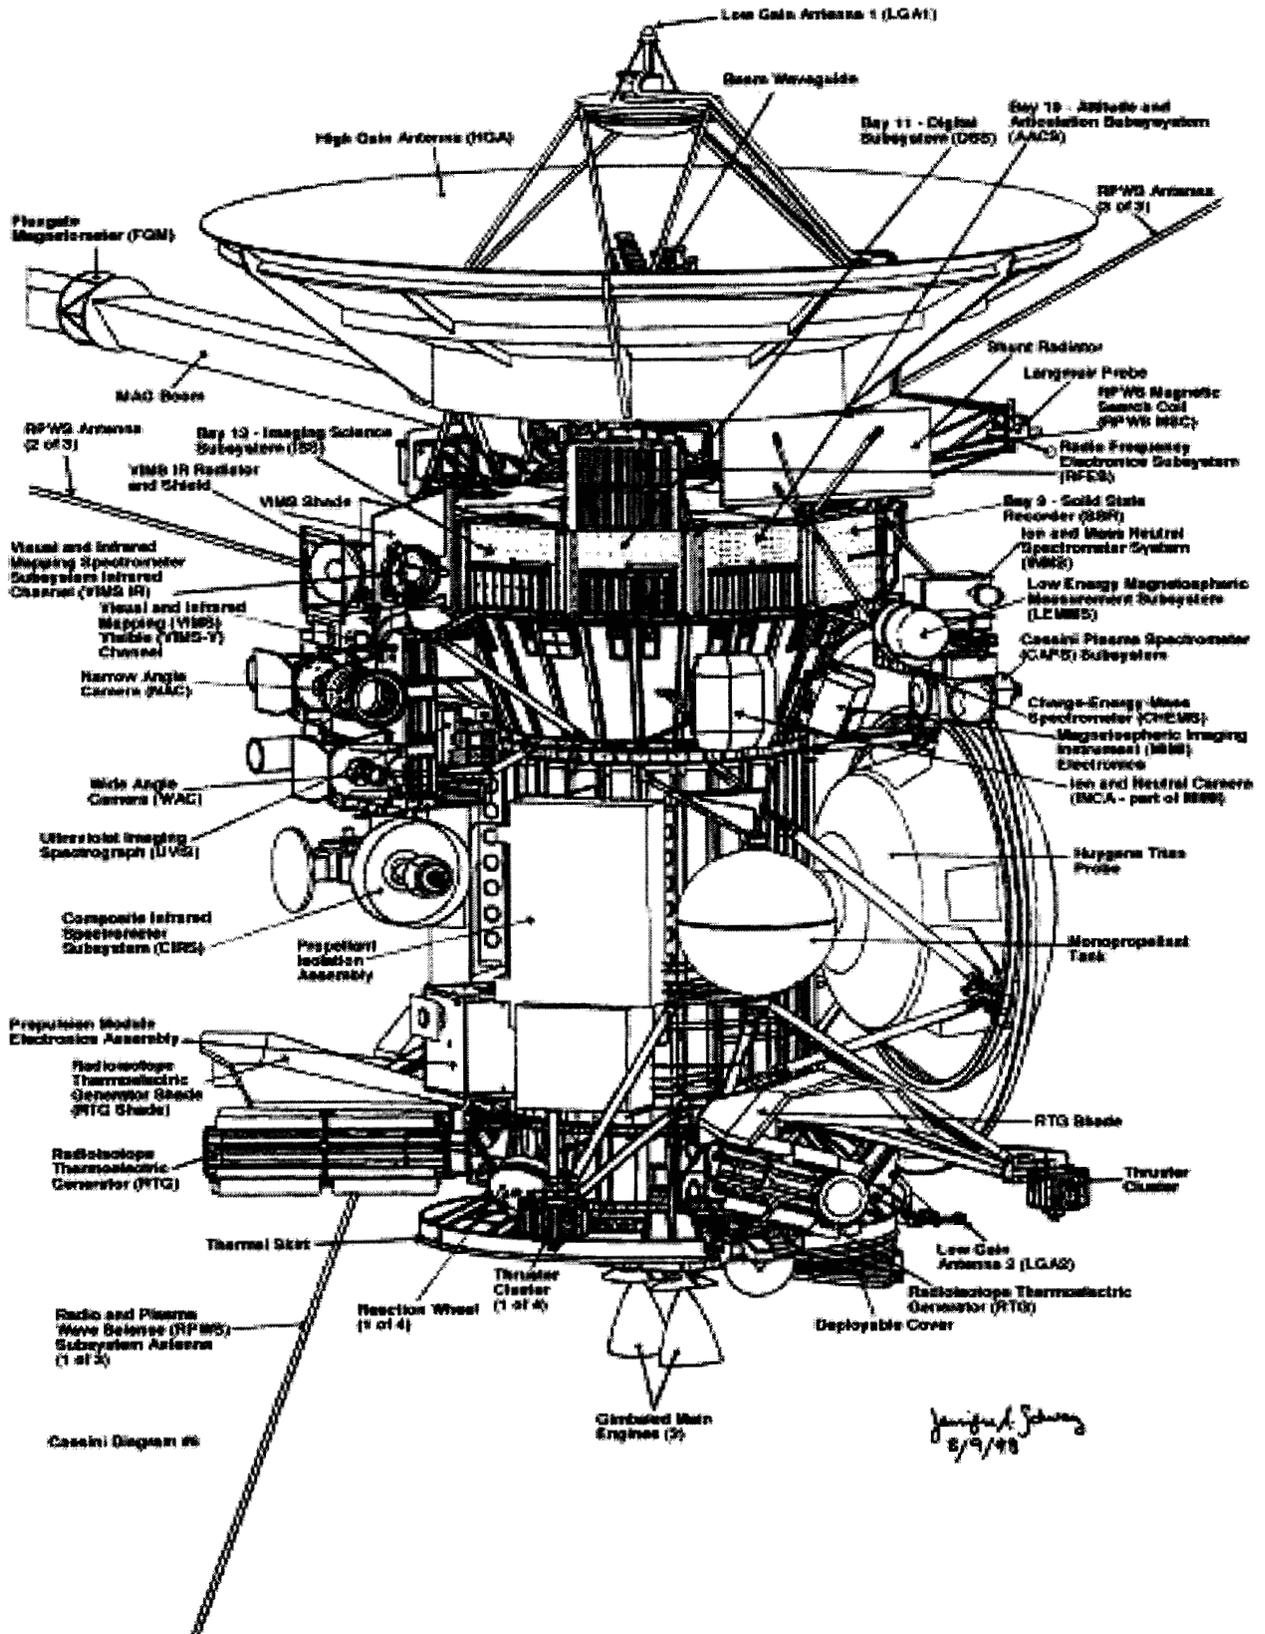


Figure 5. Cassini Spacecraft in Nominal Cruise Configuration (w/o Magnetometer Boom)

the early portion of VVEJGA; this precluded HGA communication with the Earth. The spacecraft now relies on the HGA almost exclusively, using X-band or Ka-band. S-band communications are also provided on Cassini. The orbiter primarily communicates with NASA's Deep Space Network (DSN) 34-m and 70-m stations in Goldstone, California; Madrid, Spain; and Canberra, Australia. Telecom performance has been nominal during the Cassini mission.

The Attitude and Articulation Control Subsystem (AACS) is responsible for maintaining the inertial pointing of Cassini, as well as providing control authority during spacecraft turns and spacecraft Trajectory Correction Maneuvers (TCMs). The Cassini AACS is fully redundant, and inertial knowledge is typically obtained via celestial reference or gyroscope-based estimates. AACS controls spacecraft attitude by using Reaction Control System (RCS) thrusters or reaction wheels, typically. Attitude changes are typically accomplished by firing two of eight 0.9-N hydrazine thrusters, part of Cassini's monopropellant propulsion subsystem, a part of the Cassini Propulsion Module Subsystem (PMS). In addition, AACS also controls the bipropellant portion of the PMS, including maintaining pitch and yaw control during main engine TCMs through engine gimbaling. AACS performance has been excellent throughout over four years of mission operations.

Extensive on-board fault protection against a multitude of fault conditions is provided on Cassini. These fault protection algorithms are necessarily autonomous, due to long (up to ninety-minute) one-way light times, the high demand for DSN tracking coverage (resulting in no tracking for days at a time), and the communication losses experienced around solar conjunctions. Cassini has entered fault protection or "safe mode" three times since launch, due to minor violations of spacecraft constraints. Spacecraft performance during these safing events and the subsequent recoveries was as expected.

The orbiter contains a dozen science instruments, fixed to the body of the spacecraft. This is unlike the Voyager and Galileo spacecraft, which had scan platforms for remote sensing science instruments. Therefore, the entire Cassini spacecraft must change attitude for remote sensing observations. Science instruments on the orbiter include infrared, ultraviolet, and visible light cameras and spectrometers (optical remote sensing instruments); radar and radio science subsystems (microwave remote sensing instruments); and magnetometer, plasma wave, and dust detector analyzers and spectrometers (fields, particles, and waves instruments). This suite of science instruments

will allow a thorough, cross-disciplinary investigation of the Saturnian system.

### III. PMS Hardware Summary

The Cassini Propulsion Module Subsystem (PMS) is actually comprised of two complete, separate propulsion systems, and is by far the most complex interplanetary propulsion system ever built. The PMS was built by Lockheed-Martin Astronautics under JPL/NASA contract. The Cassini PMS line drawing and propulsion schematic are reproduced in Figures 6 and 7, respectively.

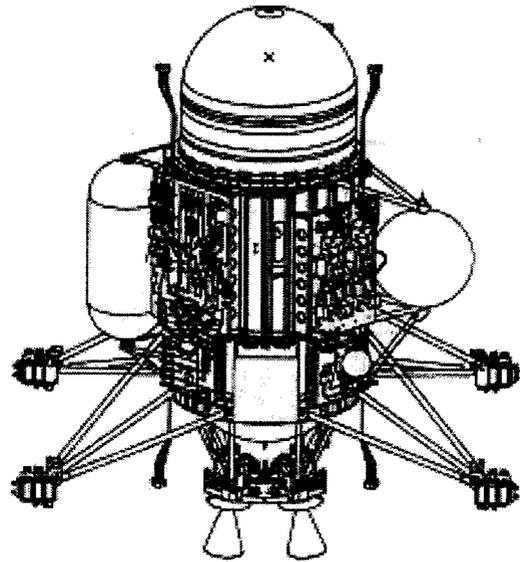


Figure 6. Cassini PMS Line Drawing

The monopropellant portion of the PMS consists of a single blowdown hydrazine tank, eight prime and eight back-up 0.9-N Voyager-heritage hydrazine thrusters, and a pyro-isolated one-time helium recharge tank. The monopropellant propulsion system, or Reaction Control System (RCS) design has rich heritage from the Voyager program. It is used for spacecraft attitude control, momentum wheel unloads, science observations with coarse pointing requirements, spacecraft turns, and RCS TCMs. The RCS also includes pressure and temperature transducers, as well as pressure-relieving liquid latch valves in the prime and back-up thruster branches. Many RCS components are located on the second Propellant Isolation Assembly (PIA-2) panel, which is evident in Figure 6. The RCS schematic is contained within the right half of Figure 7.

The bipropellant portion of the PMS is a very complex, pressure-regulated system with many redundant components. A single high-pressure,

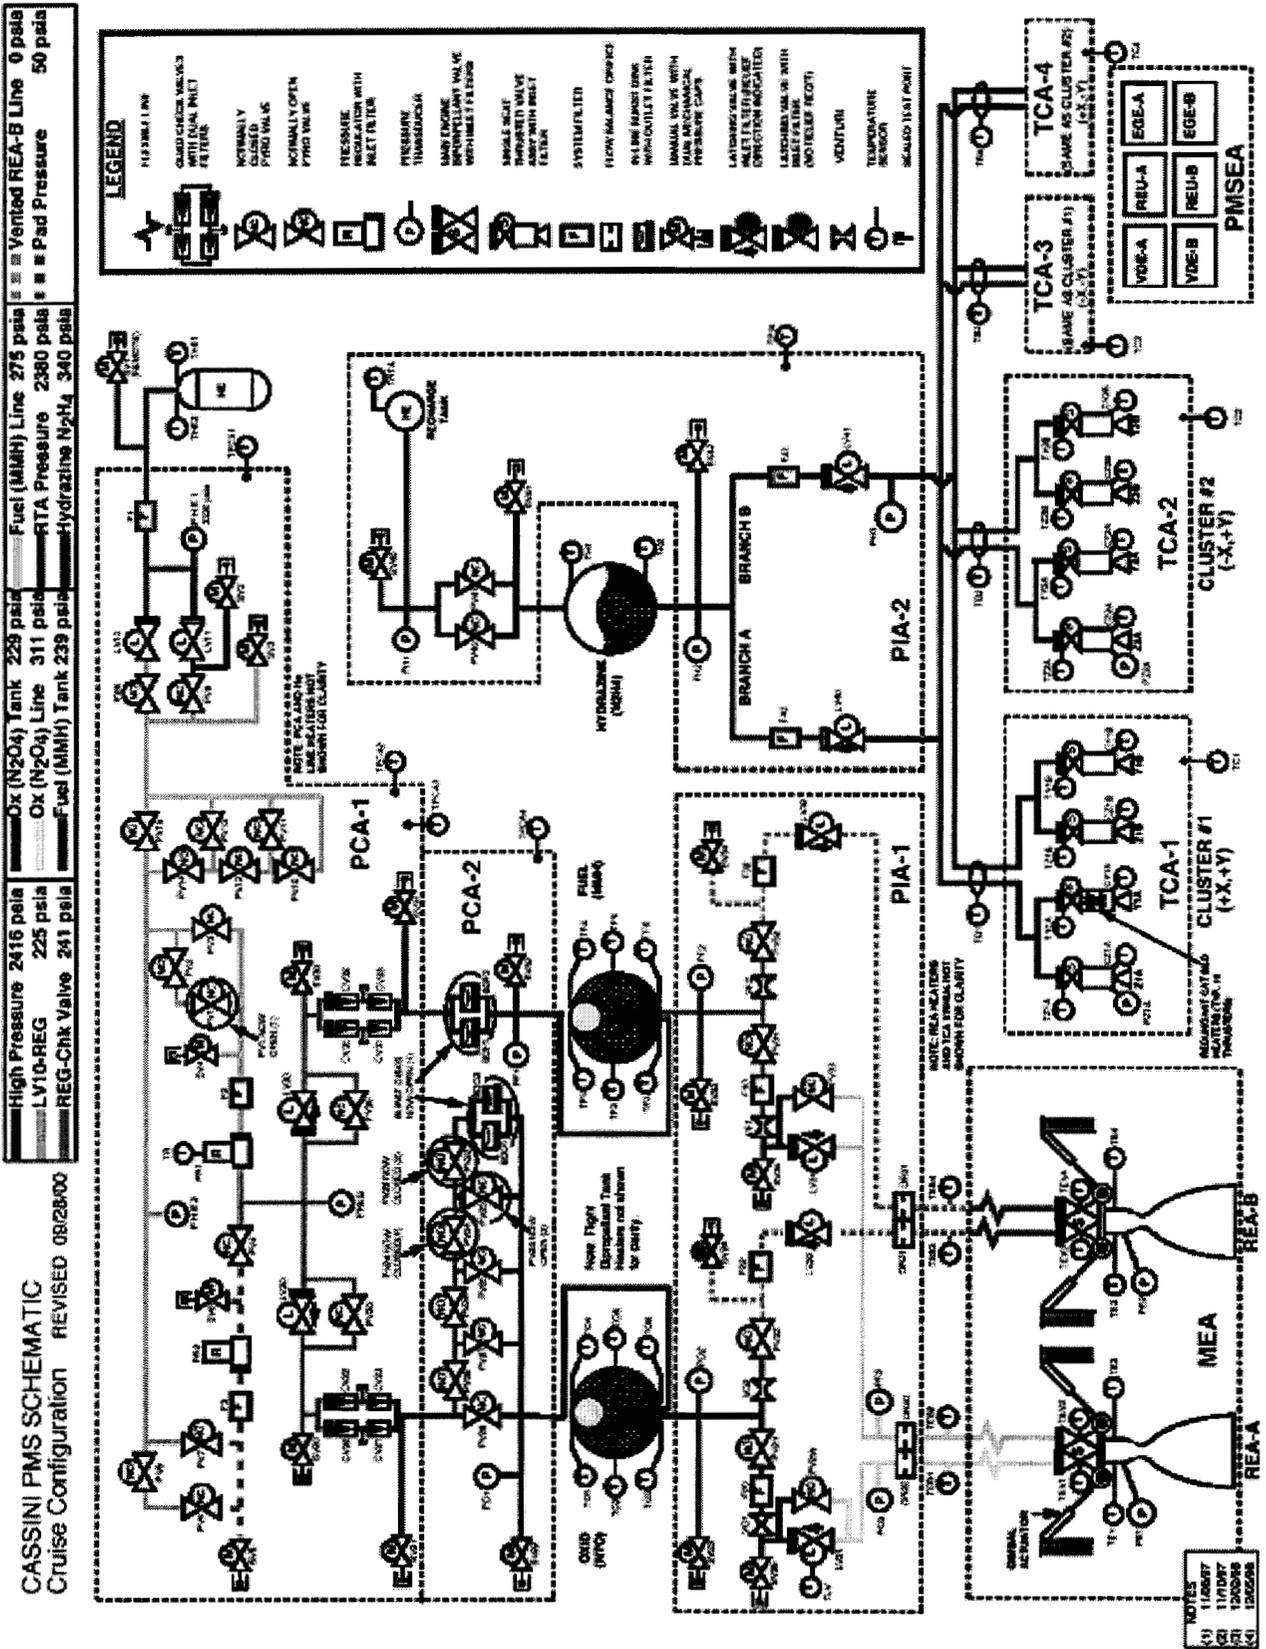


Figure 7. Cassini PMS Schematic

composite-overwrapped helium tank provides high-pressure gas during regulated maneuvers. Helium flows through the prime high-pressure latch valve (LV10) and through a high-pressure pyro-isolation ladder (PV10-PV15) to the prime pressure regulator, PR1. Note that Figure 7 includes a back-up high-pressure latch valve (LV11), as well as a pyro-isolated, pristine back-up regulator (PR2).

Downstream of PR1, the helium flow divides into two paths, through low-pressure helium latch valves (LV20 and LV30), quad-redundant oxidizer and fuel check valve packs, and an oxidizer-side pyro-isolation ladder (PV22-PV29). These features were added in the wake of the Mars Observer failure investigation, which determined that oxidizer vapor migration, condensation, and energetic reaction with fuel could have contributed to the loss of the spacecraft. One-time-use burst disks isolate the pressurization system from the single oxidizer and single fuel tank before initial pressurization. The Cassini bipropellant oxidizer is nitrogen tetroxide (NTO) with three percent mixed oxides of nitrogen (3% nitric oxide, MON-3). For the bipropellant fuel, Cassini utilizes monomethylhydrazine (MMH); NTO and MMH are hypergolic and have rich space heritage (Galileo, space shuttle RCS, etc.).

Downstream of the NTO and MMH tanks, propellant lines lead to two Kaiser Marquardt (now GD-OTS) R-4D 445-N main engines, one prime and one back-up. The second main engine was added because a single main engine was essentially the only single point failure in the original PMS design. Each engine is gimballed for pitch and yaw control during main engine burns; roll control is provided by RCS thrusters. The prime main engine, or Rocket Engine Assembly-A (REA-A), was primed after launch. This involved bringing NTO and MMH to the (closed) main engine valves by opening liquid latch valves LV21 and LV31, respectively. The back-up main engine, REA-B, remains unwetted to preclude flow decay concerns during years of inactivity during mission operations. Both REA-A and REA-B vented helium (pad pressure) to space by the opening of main engine valves on both engines. This occurred before the priming of REA-A.

Details on the mission requirements, design, and pre-launch performance of the PMS have been published.<sup>3</sup> Table 2 lists the manufacturer and the experience with the various components of the PMS; this table is reproduced exactly from the pre-launch reference.

#### IV. PMS Consumable Summary

The primary PMS consumables are monopropellant (hydrazine), bipropellant (NTO and MMH), and RCS thruster valve cycles. The usable propellant remaining is probably the most critical spacecraft consumable since it is likely to be the life-limiting resource for the mission. Latch valve cycles are tracked as well, but realistically, they will never approach consumable limits.

The hydrazine mass remaining between launch and End-of-Year (EOY) 2001 is displayed in Figure 8. There are two methods that can be used to estimate hydrazine mass remaining. First, a hydrazine consumption model attempts to "bean-count" every drop of hydrazine leaving the spacecraft. This model is of limited use due to the inaccuracies of the pulse-mode consumption model for the 0.9-N thruster. An alternative method is to simply calculate the hydrazine mass remaining from telemetered tank pressure and temperature using a thermodynamic model. This method should offer the greatest accuracy over the long term, as long as pressure transducers are not drifting (this will be verified below). The data of Figure 8 were generated using the thermodynamic model.

About 7% of the launch load of hydrazine has been used to date. The spacecraft activities that consumed the most hydrazine are labeled in Figure 8. The two RCS TCMs to date, TCM-2 and TCM-7, each used about 0.6 kg of hydrazine. The first Instrument Check-Out activity (ICO-1) used over a kilogram of monopropellant during January, 1999. This calibration period involved many spacecraft turns and fine RCS deadbands. Mosaics in RCS control particularly consumed large amounts of propellant. This is also evident in Figure 8 at approximately mission day 1150, during the Jupiter science campaign. Cassini spacecraft control during this period was typically Reaction Wheel Assembly (RWA) control, but an RWA anomaly caused the spacecraft to revert to RCS control. Science activities clocked out on thrusters, and this led to the large usage seen in Figure 8.

Other modifications to the hydrazine consumption rate are evident from Figure 8. Following Earth Swing-By (ESB) on mission day 672, the spacecraft hydrazine usage rate decreased to typical levels for RCS (20,20,20)-mrad deadbands on the (x,y,z) axes, respectively. Around mission day 815, Cassini began using the HGA exclusively, which required a finer pointing control deadband

Table 2. PMS Hardware Component Summary

CPMS Component	Supplier	Flight Heritage / Similarity
Bipropellant Tanks (BTA)	Lockheed Martin Astronautics	New Design (Qualified)
Monopropellant Tanks (MTA)	Pressure Systems Inc. (PSI)	Shuttle APU, Magellan
Helium Tank Assembly (HTA)	Lincoln Composites	New Design (Qualified)
Recharge Tank Assembly (RTA)	Arde	New Design (Qualified)
Rocket Engine Assembly (REA)	Kaiser Marquardt	IABS, Mars Observer (Qualified for expanded operational envelope)
Main Engine Assembly (MEA)	Lockheed Martin Astronautics	New Design (Qualified)
Thrusters	Olin Aerospace Corporation (OAC)	Voyager (1-N thruster with Moog valve)
High Pressure Latch Valve	Eaton	EURECA, Intelsat VI, COBE, DSP (Delta Qualified)
Low Pressure Latch Valve	Eaton	EURECA, Intelsat VI, COBE, DSP (Delta Qualified)
Propellant Biprop Latch Valve (Ti)	Vacco	Numerous commercial/military spacecraft (Delta Qualified)
Filters	Vacco	Numerous commercial/military spacecraft (Delta Qualified)
Pyro Valves (all SS)	OEA	Numerous spacecraft starting with Viking (Requalified)
Service Valves (SS & Ti)	OEA	Numerous spacecraft
Pressure Regulator	Mu Space Components	Heritage design (Requalified)
Check Valve Quad Package	Sterer	Heritage design from Galileo – requalified as individual valve and quad assembly
Flexline (Titanium)	Avica	New Design (Qualified)
Venturi (Titanium)	Flow Systems	New Design (Qualified)
Burst Disc Assembly	Hydrodyne	Heritage Design (New Supplier – Requalified)
Temperature Sensors	Rosemount	Numerous spacecraft
Pressure Transducers	Gulton-Statham	ACTS, DMSP (requalified for design modifications)
Heaters (tanks, plates, engine valves)	Tayco	Numerous spacecraft

setting of (2,2,20) mrad. Note the increased hydrazine consumption rate due to this deadband tightening. Propellant consumption was also minimized during RWA operations as compared to RCS operations, as expected. However, friction anomalies in the RWAs suggested that Cassini should be controlled in RCS mode for overall mission risk mitigation. Finally, note the improvement to the hydrazine consumption rate in RCS mode around mission day 1440. Two AACS parameters were updated to allow improved RCS control at (2,2,20)-

mrad deadbands. These parameter updates reduced thruster double pulsing at deadband firings by 6% and hydrazine consumption by a significant 41%.

A similar plot to Figure 8 was generated for the NTO and MMH mass histories, and is presented in Figure 9. Note that the vertical scale of Figure 9 is expanded for increased resolution. The thicker curve in Figure 9 represents the reconstructed NTO mass vs. mission time between launch and EOY 2001.

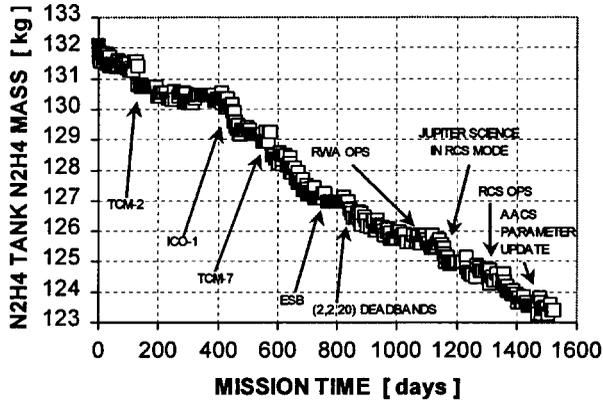


Figure 8. Cassini Hydrazine Mass vs. Time

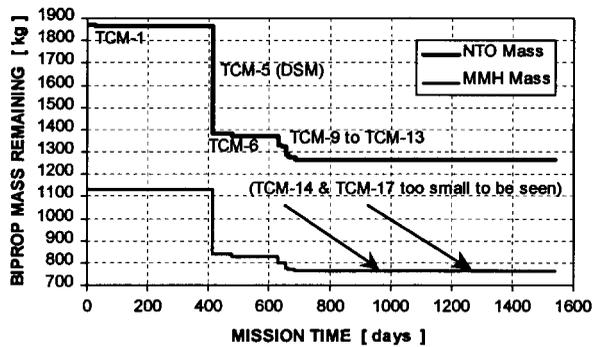


Figure 9. Cassini NTO & MMH Mass vs. Time

MMH mass is represented by the other curve in the figure. Notice that TCM-1 used an imperceptible amount of propellant (on this scale). TCM-14 and TCM-17 were even smaller main engine maneuvers; as such, their propellant usage is not discernible in Figure 9. The largest maneuver by far was TCM-5, the Deep Space Maneuver (DSM). This plane-change maneuver was nearly ninety minutes in duration, and was a fully pressurized, main-engine burn. It consumed 25% of the total bipropellant load. Bipropellant usage has generally been as expected for any given maneuver, and less than expected for the mission to date. This is due to excellent navigation, including trajectory re-optimization and the cancellation of some main-engine TCMs.

The Cassini PMS consumable summary through EOY 2001 is presented in Table 3. About one-third of the bipropellant load has been used to date. Bipropellant margin for the nominal mission

Table 3. PMS Consumables as of EOY 2001

PMS Consumable	Used thru EOY 2001	Lifetime Limit	% Used
NTO (Oxidizer) Mass [kg]	602.8	1869.0	32.3%
MMH (Fuel) Mass [kg]	365.9	1131.0	32.4%
N <sub>2</sub> H <sub>4</sub> (Monoprop.) Mass [kg]	9.4	132.0	7.1%
Z1A Thruster Valve Cycles [ ]	33460	273000	12.3%
Z2A Thruster Valve Cycles [ ]	30124	273000	11.0%
Z3A Thruster Valve Cycles [ ]	20451	273000	7.5%
Z4A Thruster Valve Cycles [ ]	24108	273000	8.8%
Y1A/Y3A Thruster Valve Cycles [ ]	11261	273000	4.1%
Y2A/Y4A Thruster Valve Cycles [ ]	7010	273000	2.6%
Z1A Thruster Throughput [kg]	1.97	25.0	7.9%
Z2A Thruster Throughput [kg]	1.94	25.0	7.8%
Z3A Thruster Throughput [kg]	1.48	25.0	5.9%
Z4A Thruster Throughput [kg]	1.51	25.0	6.0%
Y1A/Y3A Thruster Throughput [kg]	0.66	25.0	2.6%
Y2A/Y4A Thruster Throughput [kg]	0.58	25.0	2.3%
Latch Valve LV10 Cycles [ ]	5	5000	0.1%
Latch Valve LV20 Cycles [ ]	2	5000	<0.1%
Latch Valve LV30 Cycles [ ]	4	5000	<0.1%
Latch Valve LV21 Cycles [ ]	12	5000	0.2%
Latch Valve LV31 Cycles [ ]	12	5000	0.2%
Latch Valve LV22 Cycles [ ]	1	5000	<0.1%
Latch Valve LV32 Cycles [ ]	1	5000	<0.1%
Latch Valve LV40 Cycles [ ]	1	6600	<0.1%

remains positive. The margin has decreased, however, due to the Huygens probe telecommunications anomaly. This anomaly required a change to the Cassini orbital tour, including the addition of extra Titan orbits before probe release and an increased Orbiter Deflection Maneuver (ODM)  $\Delta V$ . Hydrazine mass margin remains excellent, and is currently over 40 kg for the end of a two-year extended mission following the nominal mission completion in July, 2008.

RCS thruster valve cycles are also listed in Table 3. There was some concern that valve cycle limits would be exceeded during the Cassini mission, based on early assessments of in-flight telemetry. However, as more data have been gathered at higher solar distances, the projections now show a factor of two margin against the valve thruster cycle lifetime of 273,000 pulses per thruster. Z-thrusters have experienced the most cycles in flight to date, as expected. The most used thruster, Z1A, has only used about 12.3% of its thruster valve cycle lifetime. These thrusters are essentially identical to the thrusters flown on the Voyager and Deep Space One (DS1) spacecraft. No Voyager thruster showed any anomalous behavior until about 400,000 cycles had been accumulated in flight. A few of the DS1 RCS thrusters had over 300,000 pulses on them at the time of mission completion in December, 2001. No anomalies were noted on the DS1 RCS thrusters

during over three years of mission operations. The consumable limit of 273,000 pulses in Table 3 was selected to be conservative; Voyager and DS1 flight data suggest that the limit is sufficiently conservative.

In addition to the thruster valve cycle limit, there is also a per-thruster throughput limit of 25 kg. The data of Table 3 suggest that consumable usage across thrusters is fairly even, at least within the family of Z-thrusters and Y-thrusters. The Z-thruster throughput has typically been between 1.5-2.0 kg/thruster, or 6-8% of the 25-kg consumable limit. Y-thruster throughput is quite a bit smaller, 0.58-0.66kg/thruster, or 2.3-2.6% of 25 kg. The probability of a Cassini thruster exceeding the 25 kg throughput limit is very low, at least barring an independent anomaly.

Latch valve cycles are tracked in Table 3 for completeness. However, it is obvious that latch valve cycles are nowhere near their consumable limit of either 5000 or 6600 latch valve cycles. Realistically, latch valve cycles do not even need to be tracked in the future, at least with respect to spacecraft consumables. The same may be said for B-branch RCS thrusters, since they are not planned to be used in flight.

## V. TCM Performance

Cassini has executed ten main engine (ME) TCMs and two RCS TCMs to date. All maneuvers executed successfully, within specifications. Cassini navigation has been excellent; this allowed the cancellation of TCM-3, TCM-4, TCM-8, TCM-15, and TCM-16. Table 4 is a TCM summary table for the Cassini mission between launch and EOY 2001.

TCM-1 was executed on November 9, 1997, to clean up errors due to launch and Centaur-injection dispersions. The maneuver was intended to be a regulated ME TCM, but a leaking regulator prevented TCM-1 from being executed in regulated mode. This will be discussed below in more detail. TCM-1 validated main engine function and performance following the initial bipropellant system pressurization, the REA-A and REA-B venting, and the REA-A priming.

Main engine firings like TCM-1 use accelerometer control for accurate burn termination. The accelerometer minimum burn timer setting was selected to be one second, an AACs-limit for minimum ME firing time. The maximum burn time was set 10% higher than the expected burn time, to allow for an underperforming, uncalibrated main engine. TCM-1 had a nominal accelerometer shutdown, and the  $\Delta V$  error was only 1.7%. This error was primarily due to using an accelerometer not yet calibrated in flight. The TCM-1 pointing error was 0.63°, well within specifications.

TCM-2 was performed on February 25, 1998, about sixty days prior to the Venus-1 flyby. It was a small (0.18 m/s) RCS maneuver. TCM-2 was designed to correct the execution errors of TCM-1 as well as to validate RCS TCM performance. The maneuver performance was typically as expected, except for two discrepancies, which will be elucidated below.

The TCM-2 duty cycles were not as expected, though this had little impact on the delivered  $\Delta V$ . Specifically, the Z3A thruster duty

Table 4. Cassini TCM Summary Table

T C M	TCM DATE	Man. Type	$\Delta V$ (m/s)	Mag. Err. (%)	Pt. Err. (°)	Primary Purpose of TCM
1	11/9/97	ME	2.7	1.70	0.63	Correct launch dispersions; validate ME TCM
2	2/25/98	RCS	0.179	-3.51	0.51	Correct TCM-1 dispersions; validate RCS TCM
5	12/3/98	ME	450	0.05	0.89	Plane-change maneuver for V2; validate SOI
6	2/4/99	ME	11.552	0.19	0.07	Correct TCM-5 dispersions
7	5/18/99	RCS	0.225	-2.36	1.34	Correct TCM-6 dispersions
9	7/6/99	ME	43.485	-0.14	0.12	Correct V2 flyby dispersions; Earth "walk-in"
10	7/19/99	ME	5.130	-0.05	0.11	"Walk-In" maneuver for Earth flyby
11	8/2/99	ME	36.288	-0.06	0.11	"Walk-In" maneuver for Earth flyby
12	8/11/99	ME	12.246	-0.09	0.09	"Walk-In" maneuver for Earth flyby
13	8/31/99	ME	6.685	-0.16	0.31	Correct Earth flyby dispersions
14	6/14/00	ME	0.55	2.89	1.20	Flush REA-A of NTO reaction products
17	2/28/01	ME	0.51	0.13	1.20	Flush REA-A of NTO reaction products

cycle was 98%, while the Z1A duty cycle was only 80%. The Z2A and Z4A thruster duty cycles were at an intermediate value of 90%. Explanations for this duty cycle discrepancy include gross thruster misalignment, out-of-specification thruster performance, unmodeled center-of-mass (c.m.) errors, or a combination of all three effects. The Cassini c.m. was not measured before launch, and DSM ME gimbal-angle data also suggest an unmodeled c.m. error reasonably consistent with the c.m. inferred from TCM-2 duty cycle data. Therefore, it is quite possible that the errors in thruster alignment and the engine-to-engine performance variations are within specification. That is, a “believable” c.m. error may explain TCM-2’s duty cycles.

RCS TCMs do not use the accelerometer for burn termination, since it is of insufficient accuracy to be useful for these low-acceleration burns. Rather, an essentially “timed” burn is used, with extra time allowed for thruster duty cycling during the TCM. The four prime RCS Z-branch thrusters (Z1A-Z4A) apparently delivered about 3.5% less thrust than expected during TCM-2, as can be verified in Table 4. (The magnitude error column in Table 4 represents accelerometer errors, essentially, for ME burns, and thrust magnitude errors for RCS burns.) This was noted during the Voyager mission as well, and was eventually explained as a steady-state thrust decrease caused by chamber pressure roughness.

The roughness as seen on Voyager is indicative of a minor feed system instability, and Cassini apparently shows similar behavior. Figure 10 represents the Z3A chamber pressure trace during TCM-2. This thruster was chosen because it had a duty cycle of nearly 100%. This allowed for a clearer assessment of chamber pressure oscillations throughout the burn. This effect may help explain the TCM-2 apparent reduced thrust vs. ground test. This will be explored below, in the context of TCM-7 performance,

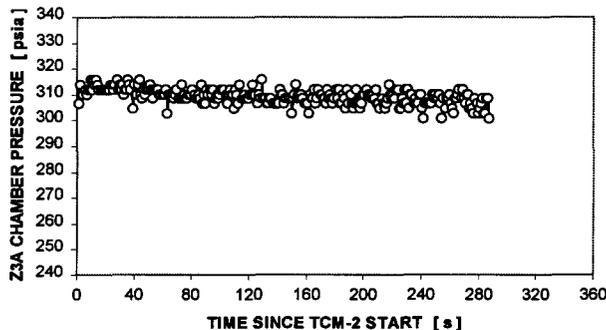


Figure 10. TCM-2 Z3A Chamber Pressure

Through excellent Cassini navigation and trajectory re-optimization, TCM-3 and TCM-4 were canceled. This was possible because DSM (TCM-5) was such a large deterministic maneuver. The first Venus flyby occurred on April 26, 1998, the first of four gravity assist flybys on the VVEJGA trajectory. Spacecraft performance was excellent throughout the flyby, and some unique science was obtained as well. DSM was executed on December 3, 1998, the first pressure-regulated ME burn in flight. The details of the pressurization system behavior during DSM will be covered below. The pre-maneuver activities included firing normally closed pyro valve PV23 to enable helium flow to the NTO tank; opening LV10 for ten minutes to repressurize the NTO and MMH tanks before DSM; and opening NTO and MMH gas-side latch valves LV20 and LV30, respectively, upstream of the check valve quad packs.

DSM was highly successful, changing the orbital plane of Cassini to set up the Venus-2 flyby conditions. The DSM burn-time was about ninety minutes, coincidentally the same burn duration as SOI. The desired DSM  $\Delta V$  was 450 m/s, and the accelerometer minimum and maximum burn times were set to one second and +5% of nominal burn duration, respectively. For non-critical maneuvers, the Cassini philosophy is always the same for minimum burn time—shut down the engine as soon as possible if there is a problem with the accelerometer. (Critical spacecraft activities are defined as activities that must be accomplished at a given absolute time to preserve the mission, such as SOI and probe relay.) The selection of the maximum burn time (and the minimum burn time for the SOI critical maneuver) is much more difficult. Experience with TCM-1 suggested that a setting of +5% offered a reasonable balance between protecting against accelerometer anomalies and PMS anomalies. DSM had a nominal accelerometer shutdown, as expected. The actual DSM burn time was 5255 seconds, only 0.73% longer than the predicted burn time of 5217 seconds. This was largely due to the very accurate prediction of the regulated pressure during this long burn.

The DSM-achieved  $\Delta V$  was virtually identical to the desired  $\Delta V$ , with an error of only 0.05%. This was largely due to in-flight calibration of the accelerometer using the data of TCM-1. The DSM pointing error was similar to the TCM-1 pointing error, slightly larger but in the same direction. Post-DSM analysis indicated an error in AACS flight software consistent with this pointing error of nearly 0.9°. This systematic error was corrected in later TCMs by allowing a small

spacecraft turn just before the initiation of the maneuver. The results from TCM-6 show that this modification had the desired effect. From Table 4, the pointing error changed from  $0.89^\circ$  for DSM to  $0.07^\circ$  for TCM-6. TCM-6 occurred on February 4, 1999, to clean up DSM execution errors. It used a maximum burn timer setting of +5%, and had a nominal accelerometer shutdown with little  $\Delta V$  error.

TCM-7 was the final RCS maneuver performed during the Cassini mission to date. It occurred May 18, 1999, and corrected TCM-6 execution errors. TCM-7 showed an additional Z1A-Z4A average thrust decrease of about 2.4%. This seems to be consistent with increased chamber pressure roughness in TCM-7 vs. TCM-2 (see Figure 11; cf. Figure 10). This explanation for RCS engine thrust decrease seems to be the most likely, because there has been no clear indication of decreased thrust during pulse-mode operation. This will be discussed further below.

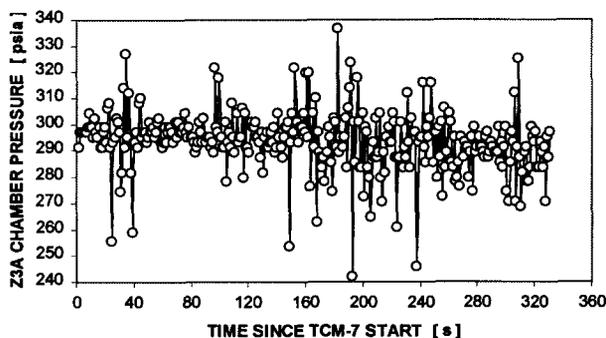


Figure 11. TCM-7 Z3A Chamber Pressure

The TCM-7 duty cycles were essentially identical to TCM-2. Again, this is consistent with an unmodeled c.m. error on Cassini. However, this had little consequence for the accuracy of TCM-7. Through excellent navigation, TCM-8 was canceled and the second Venus flyby was executed flawlessly on June 24, 1999. Again, some unique science was gleaned from this highly successful planetary flyby.

The Venus flyby conditions were selected to put the spacecraft on a near-Earth trajectory. However, in order to satisfy the maximum Earth-impact probability requirement of  $1.0e-6$ , a biased aimpoint strategy was developed for Earth approach. This led to four, main-engine, deterministic maneuvers being required in the seven weeks between the second Venus flyby and the Earth Swing-By (ESB) on August 18, 1999. TCMs 9-12 were blowdown maneuvers called “walk-in”

maneuvers, since they never directly targeted the ESB aimpoint until TCM-12.

Maneuver performance during TCMs 9-12 was excellent, as can be verified in Table 4. Magnitude errors were less than 0.15%, and pointing errors were around  $0.1^\circ$  for all maneuvers. The maximum burn time was set to +5% for all Earth-approach maneuvers, and each TCM achieved a nominal accelerometer burn termination. ESB occurred on August 18, 1999, sending the spacecraft towards the outer solar system at high velocity. Some unique lunar science observations were made during ESB. These “bonus” science opportunities offer many advantages to the Cassini flight team. They allow instrument calibration and they offer unique science observations with modern instruments. They also exercise the flight team processes and software that will be required during the orbital tour of Saturn.

TCM-13 executed on August 31, 1999, about two weeks after ESB. This allowed clean up of ESB dispersions, and maneuver performance was excellent, similar to TCMs 9-12 (see Table 4). In order to avoid potential REA-A chugging, TCM-13 was implemented as a “half-regulated” maneuver. LV10 and LV30 were opened shortly before the TCM-13 burn began; this allowed helium to flow into the MMH tank during the maneuver.

TCM-14 and TCM-17 were executed on June 14, 2000 and February 28, 2001, respectively. These two maneuvers were required to flush potential contaminants from the REA-A engine valves. Specifically, there is concern about the interaction of liquid NTO and non-Titanium components in REA-A. These chemical interactions can produce contaminants that precipitate out of NTO solution, leading to a phenomenon known as “flow decay.” The Cassini requirement is to flush each wetted engine with a minimum of a five-second firing at least every 400 days. Since REA-B has been vented but not primed, it does not require flushing.

TCM-14 and TCM-17 were designed by biasing the trajectory slightly to guarantee a burn duration of at least five seconds. TCM performance during these two TCMs was within specification, but the magnitude and pointing errors were higher than usual. This is perhaps not surprising, since the burn duration was small for these maneuvers. Incidentally, between TCM-14 and TCM-17, the spacecraft executed a highly successful and scientifically fruitful flyby of Jupiter on December 30, 2000, obtaining its fourth and final gravity assist.

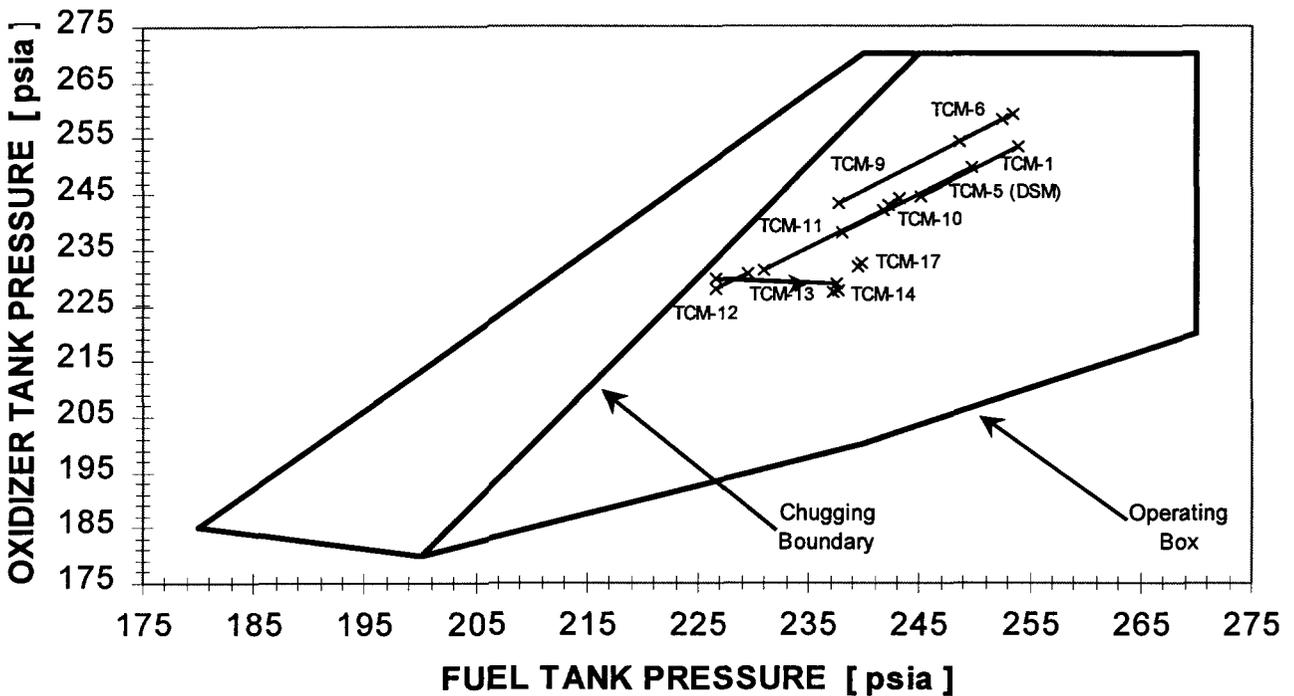


Figure 12. Cassini Main Engine TCMs/Operating Box

The Cassini main engine operating box (NTO tank pressure vs. MMH tank pressure) is presented in Figure 12, along with the flight data for the ten ME TCMs to date. The outer, hexagonal boundary of Figure 12 represents the Cassini operating box for the R-4D engine, which is more conservative than the actual allowable operating limits of the R-4D. The solid line labeled “Chugging Boundary” in Figure 12 is somewhat arbitrary. It attempts to divide the operating box into two regimes. Operation to the right of this line essentially guarantees that chugging will not occur. Operation to the left of the line suggests that chugging could start to occur at any time. Safe operation during engine chugging was demonstrated in ground test; however, the Cassini philosophy is to avoid chugging by staying to the right of the chugging boundary in Figure 12, whenever possible.

For each main engine, there is a consumable limit of 60 minutes allowed in the chugging regime. Operation for 132 minutes in the chugging regime was demonstrated during ground test. Therefore, the in-flight limit of 60 minutes represents a factor of safety of 2.2. If a burn is observed not to chug, even if it occurs to the left of the chugging boundary, it is not counted against the consumable limit.

Unobserved engine operation to the left of the chugging boundary is counted as chugging. Note from Figure 12 that TCM-13 would have crossed the chugging boundary if it were not performed as a half-regulated maneuver. The fuel-side repressurization during TCM-13 allowed an excursion towards the right in Figure 12, leading to excellent margin against chugging during Saturn approach (TCM-14, TCM-17, TCM-18, and TCM-19). This was particularly important, because TCM-13 could have been much larger than 7 m/s (see Table 4), since it was the ESB clean-up maneuver and was largely statistical.

Notice from Figure 12 that all Cassini main engine maneuvers (which have all been performed on REA-A, by the way) were executed within the operating box, as required. Moreover, all main engine TCMs occurred to the right of the chugging boundary as well. Each maneuver excursion (other than TCM-13) is shown in Figure 12 as a solid line moving from upper right to lower left. An “X” symbol marks the beginning and end of each main engine firing. All maneuvers except TCM-5 (DSM) and TCM-13 were blowdown maneuvers. DSM started with a small blowdown portion, but it was dominated by steady-state operation at the regulated

pressure (about 237 psia for both NTO and MMH tank pressure).

## VI. Pressure Transducer Drift

Many JPL missions (Voyager, TOPEX-Poseidon, and Galileo) have experienced linear pressure transducer drift during the course of their multi-year missions. This phenomenon has befuddled attempts to understand pressurization system behavior, propellant consumption, and maneuver performance during mission operations. Therefore, great pains were taken on Cassini to provide drift-free pressure transducers. This included modifications to the electronic circuitry, since this was typically implicated as the culprit for sensor drift. Specifically, an operational amplifier (op amp) in the pressure transducer supply electronics was found to drift linearly vs. time. This component instability was sufficient to explain the pressure transducer drifts seen in flight. Unfortunately, in a flight environment, there are no independent reference points for the actual pressure, so it is impossible to assess which transducers are drifting.

Pressure transducer drifts on other missions were discovered in flight by differencing the output of two, independent sensors that measure the same pressure. Therefore, missions (such as Deep Space One) that had no redundant pressure measurements were not useful for assessing drift. The maximum drift rate observed in flight was quite consistent among Galileo, Voyager, and TOPEX-Poseidon, roughly 0.24%-0.32% of full scale per year.

Figure 13 represents the difference of two Cassini NTO tank pressure measurements, PO1 and PO2, as a function of mission time between launch and EOY 2001. Note that there is an offset present at launch; in fact, this offset was well known before launch. PO2 is known to be more accurate, so the measurement from PO1 is not used in flight, except in the assessment of pressure sensor drift. The difference between PO2 and PO1 is essentially flat vs. time, suggesting that the Cassini NTO tank pressure transducers are not drifting. The slope in Figure 13 is just +0.010% of full scale per year, twenty-four times smaller than the worst-case inferred drift rate on Galileo, Voyager, and TOPEX-Poseidon. In fact, the slope in Figure 13 could easily be zero, within uncertainties.

Figure 14 is the MMH-tank analogue to Figure 13. Note that the offset between PF2 and PF1 is much smaller than for the NTO tank, typically  $1 \pm 1$  psia. As with PO1 and PO2, the fuel pressure

transducers apparently are not drifting with respect to each other. Naturally, if PO1/PO2 or PF1/PF2 are drifting at the same rate with respect to one another, this would be unobservable in Figures 13 and 14. The slope in Figure 14 is  $-0.015\%$  of full scale per year, sixteen times smaller than the Voyager inferred drift rate of 0.24% of full scale per year. Even though this slope is farther from zero than the slope of Figure 13, it still could easily be zero, within uncertainties.

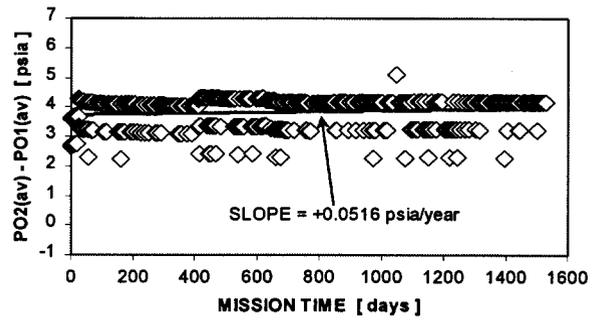


Figure 13. Cassini NTO Tank Pressure Sensor Drift

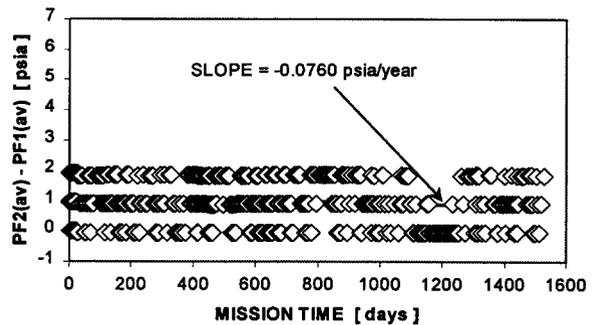


Figure 14. Cassini MMH Tank Pressure Sensor Drift

Figure 15 represents the difference between two Voyager helium tank pressure transducers vs. mission time. (To be fair, the Voyager transducers have different designs and were supplied by different manufacturers.) Note that Figures 13-15 are all on the same scale, to allow direct comparison. The data of Figure 15 are typical for a drifting pressure transducer. The slope in Figure 15 is +1.22 psia/year, which is +0.24% of full scale per year for a 0-500 psia transducer. As mentioned above, this is more than an order of magnitude larger than the Cassini "drift" rates. Galileo and TOPEX-Poseidon data look similar to Figure 15. Cassini trends will continue to be monitored during the remaining six years (at least) of mission operations. However, the data to date

certainly suggest that Cassini pressure transducer drift, if present, is minimal.

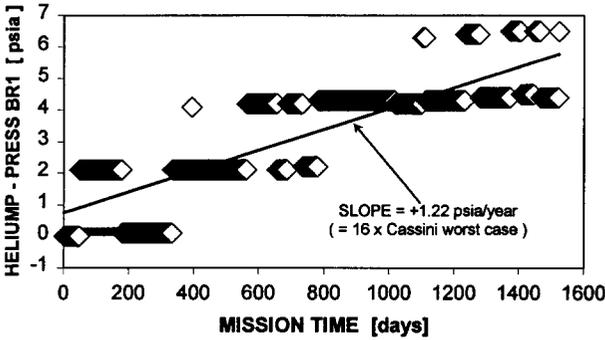
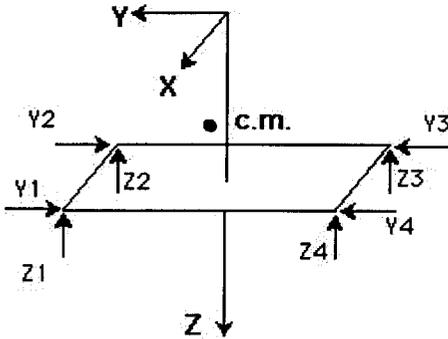


Figure 15. Voyager Pressure Sensor Drift

### VII. 0.9-N Thruster Performance

As mentioned previously, the Cassini RCS system consists of eight primary and eight back-up 0.9-N monopropellant thrusters. Only the primary, or A-branch thrusters have been used in flight. Pitch and yaw are controlled by the Z1A-Z4A thrusters, while roll is controlled using the Y1A-Y4A thrusters. The thruster geometry and axis definitions are given in Figure 16. For clarity, only one thruster is shown at each location, though each thruster location includes an A-branch (prime) and a B-branch (back-up) thruster.



Turn About Axis:	Firing Pair
+Pitch (+X)	Z3 & Z4
-Pitch (-X)	Z1 & Z2
+Yaw (+Y)	Z1 & Z4
-Yaw (-Y)	Z2 & Z3
+Roll (+Z)	Y2 & Y4
-Roll (-Z)	Y1 & Y3

Figure 16. Cassini RCS Thruster Geometry

Thruster performance during RCS TCMs was covered in a prior section, and an assessment of spacecraft turn thruster performance is incomplete and preliminary. Therefore, no RCS turn performance data are included in this work. However, Z-thruster performance during RWA unloads has been assessed during all fifty-four unloads executed to date. In addition, pulse-mode thruster performance during deadbanding (limit cycling) was analyzed in great detail. A total of 1764 pulse-mode firings were analyzed by two Cassini summer students, offering an extensive, in-flight estimation of the thruster impulse bit (IBIT) as a function of thruster on-time.

RWA unloads shed angular momentum from the reaction wheels by firing the four Z-thrusters and the four Y-thrusters. However, Y-thrusters are coupled, so a spacecraft  $\Delta V$ -based method may not be used to assess roll thruster performance. Thruster firing during RWA unloads allows the RWAs to stay within proper angular velocity ranges. The accumulation of momentum on the reaction wheels is not known a priori; therefore, RWA unloads are of variable size.

The spacecraft was largely under RWA control for nearly one year, between May of 2000 and April of 2001. During this time, a total of fifty-four RWA unloads were executed by the spacecraft. These unloads fired 125-ms thruster pulses on Z1A-Z4A, each of which imparted spacecraft  $\Delta V$ . The  $\Delta V$  was estimated by propulsion and AACs, by multiplying the number of pulses fired by the impulse bit per 125-ms pulse and dividing by the spacecraft mass. This  $\Delta V$  was calculated for each RWA unload and compared to the "actual"  $\Delta V$  as determined by navigation Doppler data. The ratio of the propulsion-estimated  $\Delta V$  to the NAV-determined  $\Delta V$  is plotted in Figure 17 for each RWA unload. Note that there is a consistent overestimation of the  $\Delta V$  by propulsion of about 5%. This agreement is quite good, which is perhaps not surprising for 125-ms pulse lengths. The data do suggest that propulsion is overestimating the impulse bit by 5%, and hence potentially the hydrazine consumption as well. RWA unloads and turns represent the best in-flight performance data for 125-ms pulses; therefore, the RWA unload data will be combined with limit cycle data below to extend the range for IBIT vs. on-time characterization.

Cassini thrusters fire pulses between 7 and 125 ms to control the spacecraft attitude in pitch, roll, and yaw during RCS cruise. Thruster on-time is determined by an AACs algorithm that attempts to

minimize propellant consumption. Before HGA operation was permanently allowed (after the Earth flyby), the RCS deadband limits were typically set to  $\pm 20$  mrad for all three axes. Solar and RTG torques tend to disturb the attitude of this three-axis controlled spacecraft. Small, paired thruster firings correct the pitch, roll, and yaw attitudes, using the thruster pairs identified in Figure 16. This operation is known as limit cycling, or deadbanding.

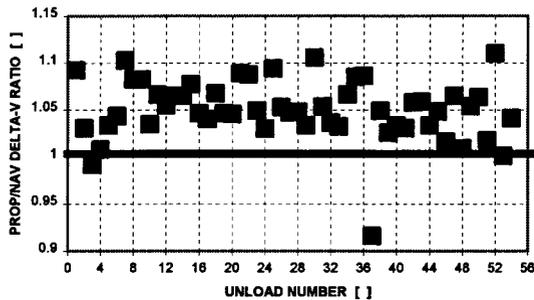


Figure 17. Cassini RWA Unload Data

Two Cassini propulsion summer students extensively analyzed thruster performance during limit cycling. Specifically, the average thruster impulse bit for a firing pair was determined as a function of the commanded thruster electrical on-time. This was possible by curve-fitting the limit cycles and estimating the angular momentum change of the spacecraft during limit cycle pulses. The spacecraft moments of inertia and thruster moment arms were required for this calculation; the Cassini propulsion team continually tracks these parameters. Moreover, the best in-flight estimate of the Cassini c.m. was used in this analysis. This estimate was based on actual gimbal positions telemetered during DSM. A total of 1764 pulse mode firings were analyzed; this represented all limit cycle firings that were amenable to analysis between launch and June, 1999.

The in-flight estimate of impulse bit vs. on-time is presented in Figure 18. No distinction is made between positive or negative pitch, roll, or yaw firings, since data consistency was generally demonstrated. Cassini limit cycle firings occurred over a wide range of thruster on-times, between 7 and 77 ms. The 125-ms points in Figure 18 were generated from the RWA unload data mentioned above. Note that the data of Figure 18 are remarkably consistent and follow the expected trend. Figure 18 includes all 1764 limit cycle firings and all 54 RWA unloads. Also note the large firing-to-firing variation for the shortest thruster on-times, also as

expected. The noise level in the impulse bit analysis is certainly expected to be higher for smaller on-times; however, the firing-to-firing variations for the smallest on-times in Figure 18 seem to be larger than expected a priori.

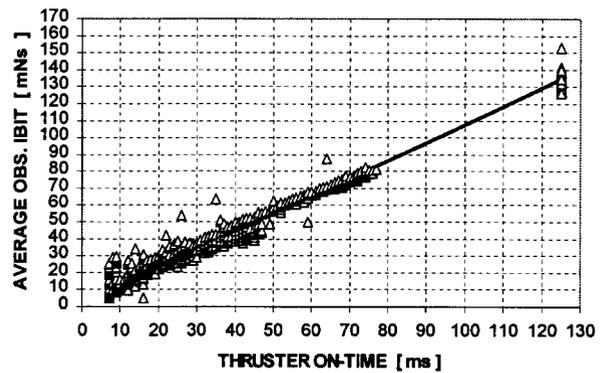


Figure 18. Cassini RCS Thruster IBIT vs. On-Time

An assessment of possible pulse-mode IBIT (or thrust) decreases vs. mission time (or throughput) was made by using the curve fit of Figure 18, eliminating thruster on-time as a variable, and plotting IBIT vs. mission time. The results were not conclusive, but they suggested that there has been no pulse-mode IBIT decreases during the Cassini mission. This is in contrast to the steady-state thrust decreases noted in TCM-2 and TCM-7, but it is consistent with the chamber pressure roughness being the culprit. This is because chamber pressure roughness is a feed system instability, which takes some time to be manifest (i.e., it is an “onset” phenomenon). As such, no pulse-mode IBIT decreases would be expected if chamber pressure roughness were the mechanism for steady-state thrust decreases.

The thorough assessment of impulse bit vs. on-time allowed a recalibration of the hydrazine propellant consumption model used on Cassini. The default hydrazine consumption model for Cassini was based on ground-test data for the 0.9-N thruster with a Wright thruster valve. Cassini utilizes a Moog thruster valve, like Voyager, so some discrepancies were expected. Using the pre-launch impulse bit and propellant consumption model, the remaining hydrazine mass overestimates the tank model (Figure 8) by about 18% between launch and EOY 2001. This can be verified in Figure 19.

In contrast, using the curve fit of Figure 18 for the impulse bit and hence hydrazine consumption model, a much closer agreement between the tank

model and consumption model is realized. This is evident in Figure 20, which is analogous to Figure 19, except now the in-flight calibration for impulse bit vs. on-time (from Figure 18) is used to calculate hydrazine consumption. Note in this case that the discrepancy between the two models decreases; in fact, the overestimate of the consumption model shrinks from 18.4% to 8.5%. The remaining 8.5% is reasonably small, and is probably due to uncertainties in pulse-mode specific impulse in the consumption model. Given the excellent transducer drift stability discussed above, the tank model probably represents the best model for remaining hydrazine mass.

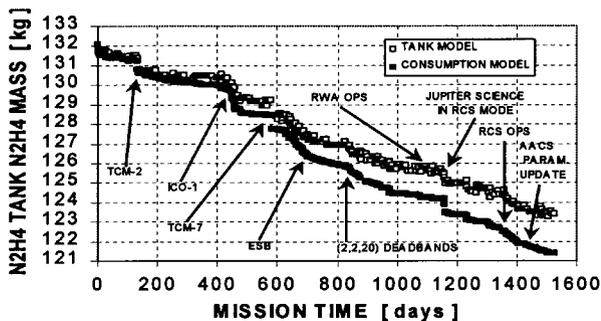


Figure 19. Comparison of  $N_2H_4$  Mass Models (1)

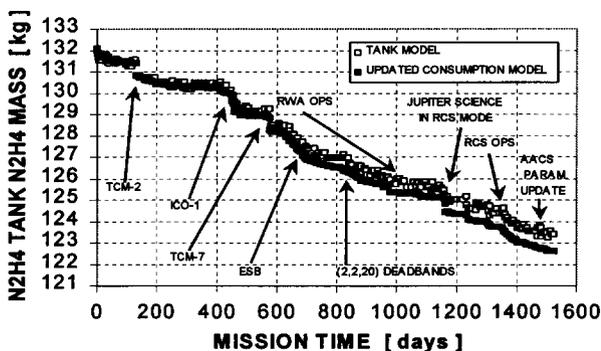


Figure 20. Comparison of  $N_2H_4$  Mass Models (2)

### VIII. Cassini PMS Helium Budgets

The Cassini propulsion team tracks the amount of helium on board the spacecraft as a function of time. This so-called “helium budget” was a useful exercise for the Galileo mission. It allowed assessment of possible internal and external helium leakage, pressure transducer drift, characterization of regulator and check valve performance, calculation of helium solubility time constants in NTO and MMH, and a crosscheck of bipropellant consumption. For

these reasons, a Cassini helium budget exercise was undertaken, for both the monopropellant and bipropellant systems.

#### Monopropellant Helium Budget

The helium budget for the hydrazine monopropellant portion of the PMS is rather different than the bipropellant helium budgets of Galileo and Cassini. Specifically, errors in hydrazine consumption (typically due to pulse-mode propellant consumption modeling uncertainties) make it very difficult to assess the helium content of the hydrazine tank. Since there is no independent way to estimate the helium content of the hydrazine tank, the total helium in the hydrazine tank is taken as constant. This assumption was required to generate Figures 19-20 above. The remaining discrepancy in Figure 20 (8.5%) may be explained by errors in pulse-mode specific impulse (most likely), hydrazine tank pressure transducer drift (less likely), or by external leaks (least likely). Initial studies indicate that the Cassini specific impulse assumptions may be overly pessimistic; this is consistent with Figure 20.

If the discrepancy in Figure 20 is attributed to hydrazine tank pressure transducer drift, the sensor drift rate would have to be three times the Galileo, Voyager, and TOPEX-Poseidon drift rates. This is deemed unlikely, though there is no way to independently verify pressure transducer drift in the hydrazine tank, with one exception. An independent pressure measurement could be obtained by opening LV41, the B-branch RCS liquid latch valve, since there is a pressure transducer downstream of LV41. Indeed, there are plans to do this for SOI and probe relay, in 2004 and 2005, respectively (to provide for fault conditions with the A-branch RCS). However, LV41 has not been opened to date, so no drift assessment is possible. An external helium leak could explain the discrepancy in Figure 20; however, there is no independent evidence of external leakage (e.g., from navigation tracking or AACS spacecraft rate data).

The pyro-isolated helium recharge tank in the hydrazine system offers a less trivial application of the helium budget method. The helium content of the recharge tank (as determined from telemetered tank pressure and temperature) is presented in Figure 21, again between launch and EOY 2001. As expected for this pyro-isolated tank, the helium content is essentially constant vs. time.

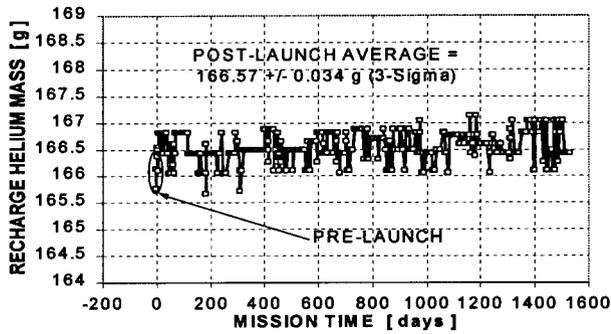


Figure 21. Recharge Tank Helium Mass vs. Time

There does appear to be a slight upward drift in Figure 21, which may be the manifestation of a slight pressure transducer drift. A linear fit of the Figure 21 data suggests a drift rate of no more than +0.015%/year of full scale. This is more than an order of magnitude smaller than the drift rate inferred for the Galileo helium tank, and could easily be zero within uncertainties. Therefore, the most likely scenario is that the recharge tank is leak-tight and has a non-drifting pressure transducer. It is also likely that PV40 and PV41 are leak tight, though the specification for a closed pyro valve leak rate is so stringent that out-of-spec leakage would not be evident in Figure 21.

Incidentally, the recharge tank helium load of  $166.57 \pm 0.034$  grams ( $3\sigma$ ) is very close to the target load of  $167.78 \pm 0.34$  grams ( $3\sigma$ ), but this implies that the target load was not met. However, the discrepancy is not large, and the hydrazine tank was slightly overloaded with helium vs. requirements, so this will have little consequence for mission operations. Sufficient helium was loaded into the monopropellant system to execute the prime and extended Cassini missions. The firing of PV40 and/or PV41 will not occur until 2005 at the earliest, well after probe relay. This will increase the hydrazine tank pressure, allowing for more thruster control authority late in the mission (during low-altitude Titan flybys, for example).

### Bipropellant Helium Budget

The Cassini propulsion team tracks the total helium mass in the bipropellant system as a function of time. This total should remain nearly constant, since only about 141 grams of helium out of approximately 8460 grams loaded is dissolved in NTO and MMH and thus expelled during ME TCMs.

Nominally, helium should just shift from the pressurant tank to the bipropellant tanks during periods of regulation (e.g., fully regulated TCMs, half-regulated TCMs, initial pressurization, pre-maneuver pressurizations, fuel-side pressurizations, etc.). Therefore, since the amount of helium in each tank is calculable from spacecraft telemetry, the total may be added up to verify that it is constant vs. time.

Figure 22 represents a typical plot of the Cassini bipropellant helium budget between launch and EOY 2001. The vertical scale does not go through zero, in order to see more detail in the data. The bottom curve in Figure 22 is the helium tank helium mass; the successively higher curves include the helium in NTO ullage, MMH ullage, NTO solution, and MMH solution, respectively. Therefore, the top curve of Figure 22 represents the total helium in the bipropellant system. Note that the helium tank helium mass data noise dominates the noise in Figure 22. This is as expected, because the largest uncertainty in the calculation of helium mass is the coarse resolution of the helium tank pressure transducer.

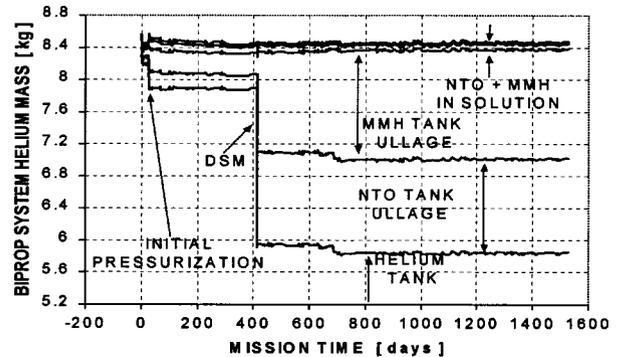


Figure 22. Cassini Bipropellant Helium Budget

From Figure 22, it can be seen that large amounts of helium were transferred during the initial pressurization and the DSM; however, as expected, the total helium mass remained essentially constant. This can be verified analytically by tabulating the helium masses transferred during pressurization events (see Table 5). Note from Table 5 that there were additional pressurization events, not labeled in Figure 22 for clarity. They were the fuel-side pressurization between TCM-9 and TCM-10 and the half-regulated TCM-13 burn. In all four cases, the amount of helium transferred as calculated by the decrease in helium tank (HTA) helium mass agreed very well with the amount calculated by the increase in bipropellant tank (BTA) helium mass. These

results are as expected, for the case of zero external helium leakage (within uncertainties).

Table 5. Cassini Helium Mass Cross-Check

EVENT	NTO $\Delta$ HE MASS [ kg ]	MMH $\Delta$ HE MASS [ kg ]	BTA $\Delta$ HE MASS [ kg ]	HTA $\Delta$ HE MASS [ kg ]	HTA/ BTA ERR. [%]
Initial Press.	0.139	0.181	0.320	-0.309	-3.4
DSM	0.947	0.937	1.884	-1.831	-2.8
Post-TCM9	0	0.362	0.362	-0.385	+6.4
TCM-13	0	0.758	0.758	-0.769	+1.5

More detail can be gleaned from the helium budget by expanding the scale of Figure 22. In Figure 23, only the top curve from Figure 22 is presented. Note that the total helium mass is initially overestimated. This was unanticipated, but eventually was explained by in-flight helium tank temperature gradients caused by waste RTG heat near the bottom of the spacecraft. These gradients were not observed in telemetry, because the two helium tank temperature measurements are both located near the top of the helium tank.

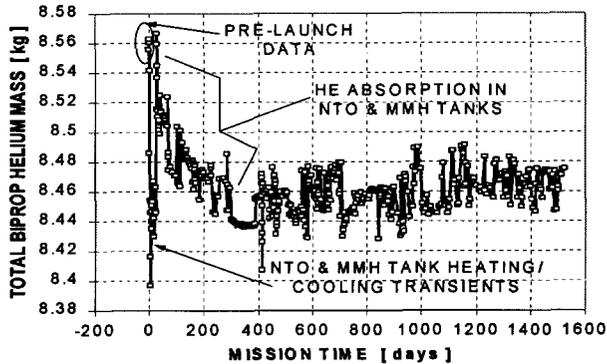


Figure 23. Cassini Total Bipropellant Helium Mass

Following launch, the NTO and MMH tanks were thermally conditioned (“ratcheted”) to avoid overpressurization during the first Venus flyby. This was important because the Cassini NTO and MMH tank ullage volumes were only 6-7% at launch. Thermal conditioning consisted of turning on 30W heaters on the NTO and MMH tanks for many days. This disturbed the thermal equilibrium in the propellant tanks, as expected. Since the helium budget calculation assumes thermal equilibrium, it is not surprising that the bipropellant tank heater transients disturb the helium budget. This is evident in Figure 23.

The data of Figure 23 were generated assuming complete saturation of helium into liquid NTO and MMH. However, helium saturation is a diffusion-limited process, and actually occurs over many days. Therefore, the helium budget as presented in Figure 23 should overestimate the total helium mass (e.g., after initial pressurization), eventually decaying to the expected level as helium is absorbed into liquid NTO and MMH. This behavior is readily apparent in Figure 23, between mission days 25 and 415, roughly.

DSM occurred on mission day 415, and the data between DSM and EOY 2001 are presented in Figure 23 as well. Note that there may be a slight upward trend in helium mass over this time span. It is possible that helium tank pressure transducer drift is producing this apparent trend. However, this was investigated solely for the helium tank data, since the errors in Figure 23 include helium mass errors from the NTO and MMH tanks as well as the helium tank.

Figure 24 shows the helium tank helium mass as a function of mission time. Note that the four pressurization events to date (initial pressurization, DSM, post-TCM9 pressurization, and TCM-13) are easily discernible in Figure 24. An assessment of possible helium tank pressure transducer drift and LV10 leakage was undertaken between pressurization events. No clear evidence was found for transducer drift, since the inferred drift rate changed significantly between pressurizations, in both magnitude and direction. In contrast, the transducer drift on other JPL missions showed a constant, linear drift rate over time. The magnitudes of the inferred drift rates from Figure 24 were typically an order of magnitude smaller than drift rates on Galileo, Voyager, and TOPEX-Poseidon. This suggests that the Cassini helium tank pressure transducer is not drifting.

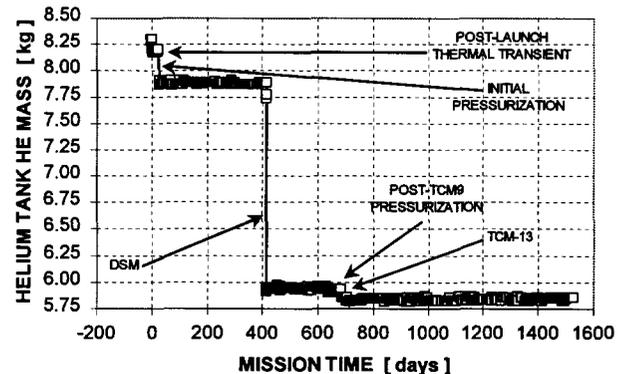


Figure 24. Helium Tank Helium Mass vs. Time

The data between initial pressurization and DSM do show a slight negative slope, consistent with minor LV10 leakage. Even if LV10 leaked between initial pressurization and DSM, however, the leak rate was about 12% of the spec leak rate leakage of 20 scch. It is impossible to separate out the effects of LV10 leakage and pressure transducer drift, but the data suggesting LV10 leaked between initial pressurization and DSM are weak at best. Moreover, there is no way to independently verify the leak-tightness of LV10 by calculating the MMH tank helium mass, because this was during the time of helium saturation in liquid MMH. The most likely scenario is that LV10 has remained leak-tight (or at least within specification) following all open/close cycles, including the nearly 2.5 years between TCM-13 and EOY 2001.

Figure 25 shows the total helium content of the NTO tank as a function of time. The post-launch thermal transient due to tank thermal conditioning is evident, as is the helium transfer to the NTO tank during initial pressurization. As before, the data of Figure 25 assume complete helium solubility, so the helium mass is initially overestimated until helium is absorbed into liquid NTO. This occurs between mission days 25 and 415. Helium absorption is driven by diffusion, so the exponential behavior in Figure 25 is anticipated. Details on helium saturation in NTO will be provided below.

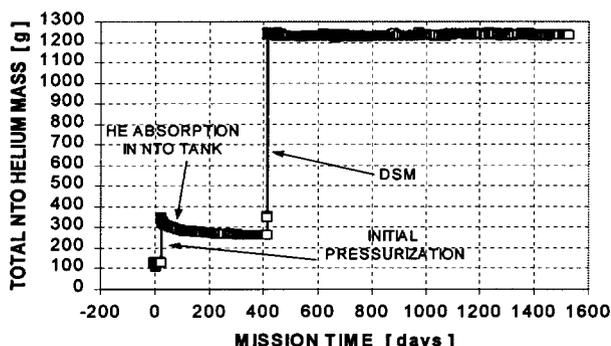


Figure 25. NTO Helium Mass vs. Time

Following initial pressurization, PV22 was fired closed. This was done to prevent NTO vapors from migrating upstream in the Cassini PMS pressurization system. Since the NTO tank was pyro-isolated between just after initial pressurization and just before DSM, there is no reasonable expectation that helium could enter the NTO tank, even given a leaking LV10. Therefore, the MMH tank helium

budget may prove more useful for assessing LV10 leakage between initial pressurization and DSM.

Just before DSM, PV23 was fired open to allow a helium path to the NTO tank. Note from Figure 25 the large transfer of helium to the NTO tank during DSM, as expected. Just after DSM, PV24 was fired closed, again isolating the NTO tank from the pressurization system. As expected, the post-DSM helium budget in the NTO tank demonstrated no pyro valve leakage.

Figure 26 is the MMH tank analogue to Figure 25. Note that the post-launch thermal transient effect on the helium budget is much smaller for the MMH tank than for the NTO tank. This is as expected, because the NTO tank has higher thermal gradients (due to RTG waste heat) and because NTO has a larger vapor pressure dependency on temperature. Note also from Figure 26 that the amount of helium absorbed into MMH following initial pressurization is much lower than the amount absorbed in NTO. This is as expected, and will be discussed further below.

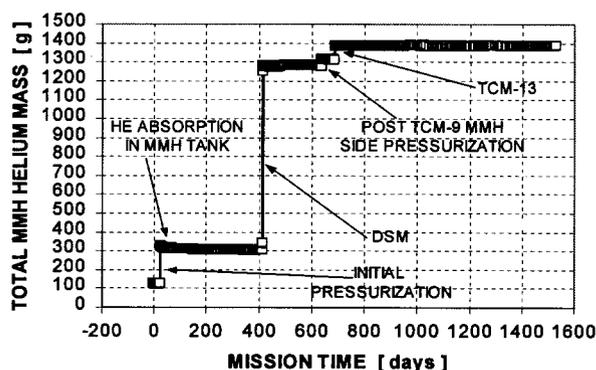


Figure 26. MMH Tank Helium Mass vs. Time

Figure 26 mimics Figure 25 except for the two additional fuel-side pressurization events after DSM. Even though the amount of helium transferred during the post-TCM9 fuel-side pressurization and TCM-13 (a half-regulated burn) is small, it clearly shows up on the scale of Figure 26. Incidentally, the data of Figures 24-26 were used to generate Table 5 (mentioned above).

Assessments of LV10 leakage were made by looking at the MMH tank helium budget between pressurization events following DSM. LV10 leakage would be apparent in the MMH tank helium budget before it would ever be discernible in the helium tank

helium budget, due to the coarse resolution of the helium tank pressure transducer. The data of Figure 26 also suggest that LV10 has not leaked during the Cassini mission, with the possible exception of the period between DSM and the post-TCM9 pressurization. Even in the most pessimistic assessment, the LV10 leak rate during this time was no more than 12% of spec, as mentioned above. Linear fits of the data of Figures 24-26 typically showed much smaller variation than the day-to-day variability in the data, so this is certainly an exercise in “noise” analysis.

Figure 27 shows the data of Figure 25 between the initial pressurization and the DSM, on a greatly expanded vertical scale. This allows a determination of the post-pressurization helium solubility into NTO. As expected, the helium mass is initially overestimated before helium is absorbed into NTO solution. Helium mass vs. time follows an exponential, as expected for this diffusion-limited process.

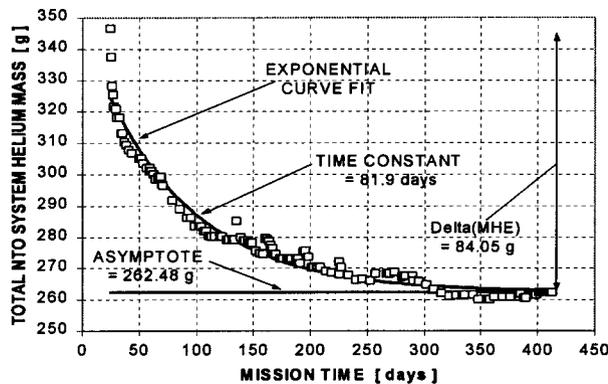


Figure 27. NTO Helium Saturation Data

An exponential fit of the data was prepared; it suggested that the time constant for helium saturation into NTO was about 81.9 days. This is a very long time, indeed, but this was expected a priori. The Galileo time constant for helium saturation into NTO was nine days, but that was for a spinning spacecraft. Mars Observer ground test data suggested a helium saturation time constant of roughly 50 days for NTO and MMH, much closer to the Cassini in-flight results. Moreover, the Cassini time constant would be expected to be longer than Mars Observer’s and Galileo’s, given the large Cassini propellant tanks and bipropellant load as compared to these two spacecraft.

The additional amount of helium absorbed into NTO solution following initial pressurization was about 84.05 grams, as can be verified in Figure 27. This agrees fairly well with the predicted value of 68.84 grams, obtained by using JPL propulsion standard solubility models. In summary, the helium saturation behavior in the Cassini NTO tank following initial pressurization was as expected, with respect to the time constant for saturation and the helium solubility level.

Figure 28 is the MMH analogue to Figure 27, and it displays very similar behavior, as expected. The time constant for helium saturation was determined by fitting an exponential curve to the in-flight data. The time constant turned out to be 80.8 days, remarkably similar to the value of 81.9 days obtained for the NTO tank. From Figure 28, the additional amount of helium absorbed into MMH solution following initial pressurization was about 23.63 grams. This is in excellent agreement with the predicted value of 25.74 grams, again from the JPL propulsion standard solubility models. In summary, the helium saturation behavior in the Cassini MMH tank following initial pressurization was as expected, and was quite similar to the behavior in the NTO tank.

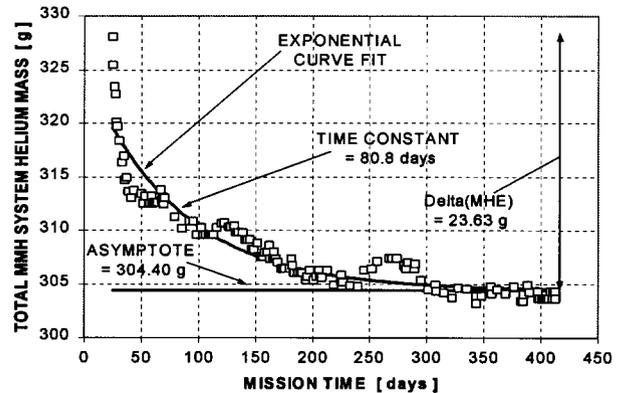


Figure 28. MMH Helium Saturation Data

## IX. Cassini Regulator Leakage

The Cassini regulator is a hard-seat device, unlike the Galileo regulator, which used a soft-seat design. The prime Galileo pressure regulator exhibited excellent leakage behavior in flight, but duplicating the Galileo soft-seat design for the Cassini regulator proved to be difficult. Therefore, the decision was made to fly Cassini with one prime and one back-up hard-seat regulator. Since hard-seat

regulators are more likely to leak due to particulate contamination, Cassini was designed to be fault tolerant in the face of the regulator leakage. Mitigating factors for Cassini include a prime and back-up high-pressure latch valve (LV10 and LV11, respectively), a pyro-isolation ladder upstream of the regulators (PV10-PV15), filters, and a pristine back-up regulator currently pyro-isolated (PR2).

Before TCM-1, the Cassini initial pressurization sequence was executed on the spacecraft. This sequence fired pyro valve PV1 to open the helium path to the propellant tanks. However, this action was performed with LV10 closed, so LV10 needed to be opened to start the flow of helium. The LV10 opening occurred as expected on November 8, 1997, and the NTO and MMH tanks quickly pressurized. However, bipropellant tank pressures continued to rise above the expected point at which the regulator was expected to lock up. In fact, the NTO and MMH tank pressures increased linearly at an astonishing rate of one DN every fifteen minutes, or 7.6 psia/hour. Since regulator leakage was deemed to be one of the more likely failure scenarios for Cassini, a real-time command to close LV10 had already been prepared and was sent to the spacecraft within a few hours of initial pressurization.

The most likely explanation for the pressurization system behavior at initial pressurization is hard-seat regulator leakage, perhaps due to particulate contamination caused by the PV1 pyro valve firing. Other scenarios were investigated as well, including a shift in the regulator set point or a scratched sapphire ball in the hard-seat regulator. These explanations were eventually dismissed as less likely, given the increase in the regulator leak rate noted at DSM (discussed below).

A determination of the initial pressurization regulator leak rate can be made by two independent methods. First, the leak rate can be estimated by measuring the NTO and MMH tank pressure rise rate and solving for the helium mass flow rate that is consistent with the pressure rise rate. This method is presented in Figure 29, and it suggests an enormous regulator leak rate of 1636 sccm, or 98160 scch. The specification for leakage for this regulator is 36 scch, so the apparent regulator leak at initial pressurization is more than three orders of magnitude larger than spec. The results of regulator ground tests before launch typically showed no discernible leakage, with one test showing leakage only slightly above specification. Clearly, this in-flight behavior represented an entirely different regime.

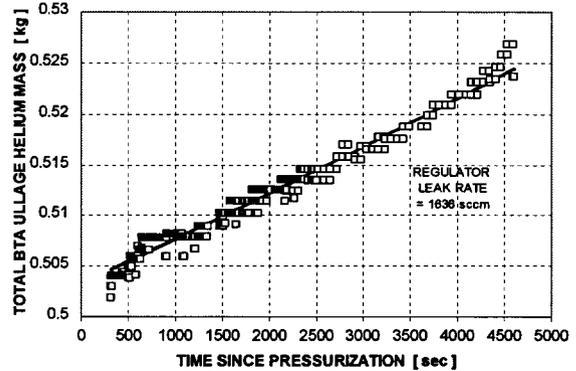


Figure 29. BTA He Mass During Init. Pressurization

A completely independent method can be used for cross-checking the regulator leak rate inferred from Figure 29. Following LV10 closure, the helium upstream of the regulator but downstream of LV10 bleeds through the leaking regulator into the bipropellant tanks. This causes the pressure between LV10 and the regulator to decay exponentially, with a pressure decay slope that depends on regulator leak rate. Figure 30 displays PHE2, the pressure between LV10 and the regulator, just after the LV10 closure. The regulator leak rate at the time of the closure of LV10 is directly proportional to the initial slope of the exponential curve fit in Figure 30. The apparent regulator leak rate from Figure 30 turned out to be 1667 sccm, or 100020 scch. This is only 1.9% higher than the leak rate determined from Figure 29, demonstrating the consistency of the two methods for calculating leak rate.

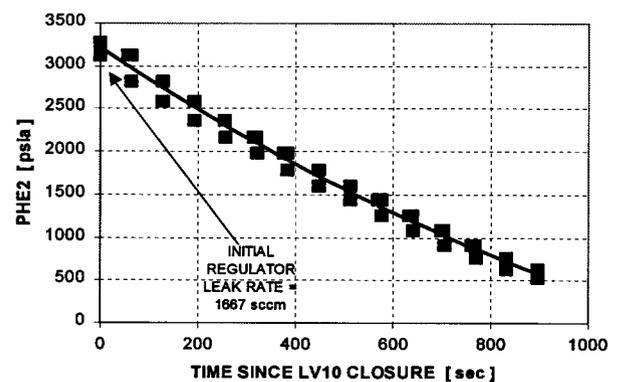


Figure 30. PHE2 Decay Curve at Init. Pressurization

Possible causes for the apparent regulator leak were investigated after TCM-1. An upward shift of the regulator lock-up pressure could explain the

behavior seen in Figure 29; however, there was no indication that the pressure rise rate was starting to decrease by the time of LV10 closure. If a regulator lock-up shift is postulated, it must have been severe enough that the regulator was far from its new lock-up point. There is no known mechanism for such a large shift in regulator lock-up point, so this explanation was deemed unlikely.

Particulate contamination can quite easily explain the regulator leak rate inferred from Figures 29-30. In fact, a trapped particle with a diameter no larger than 0.18 micron could cause a regulator leakage of 100000 scch at initial pressurization. Such a particle is much smaller than the particle size that would be stopped by the Cassini 15-micron filters. Firing debris from PV1 or residual ground contaminants are possible sources for a 0.18-micron particle. In summary, the Cassini regulator leak at initial pressurization was most likely caused by particulate contamination of the hard-seat regulator.

The next opportunity to assess possible regulator leakage came a few days before the DSM, in early December, 1999. LV10 was opened for ten minutes a day before DSM, to repressurize the NTO and MMH tanks for the maneuver. It was readily apparent that the prime regulator was still leaking; in fact, it was fairly obvious that the leak rate increased substantially since initial pressurization. Figure 31 is the pre-DSM analogue to Figure 29, representing the NTO and MMH total helium mass (in the ullage volumes) during the leaking regime of the regulator.

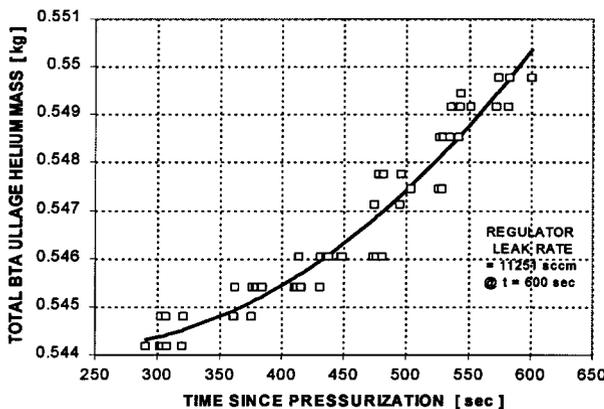


Figure 31. BTA He Mass During Pre-DSM Press.

A second-order fit of the data in Figure 31 was made, to better capture the slope at the time of LV10 closure. The reason for the non-linear behavior in Figure 31 is not yet understood, and it

will be investigated further. The marked increase in regulator leak rate is evident in Figure 31. The pre-DSM regulator leak rate was determined to be 11251 scch (675060 scch) at the time of LV10 closure. This is 6.6 times larger than the apparent leak rate during initial pressurization, suggesting that an even larger particle became trapped in the hard-seat regulator. However, particle size scales linearly with leakage rate, so a particle size of only 1.2 microns is required to explain the pre-DSM leak rate. This is still an order of magnitude smaller than the filter rating of 15 microns. Note well that the required regulator stroke during main engine firing is only 7-8 microns, so the Cassini regulator “leak” rate quoted here might be better characterized as a “flow” rate. An assessment was undertaken of SOI performance with a 15-micron particle trapped in the regulator. Though tank pressures increase throughout the ninety-minute burn in this scenario, they only violate the main engine operating box by one psia.

Figure 32 is the pre-DSM analogue for Figure 30, again offering a cross-check on regulator leak rate by looking at the upstream pressure decay following LV10 closure. The regulator leak rate from Figure 32 was estimated at 11113 scch (666780 scch), only 1.2% lower than the estimate from Figure 31. This consistency justifies the use of a second-order curve fit in Figure 31, though again there is no mathematical justification for doing so. Note that it is very difficult to discern the exponential curve fit from the data of Figure 31, suggesting that this physical model is appropriate.

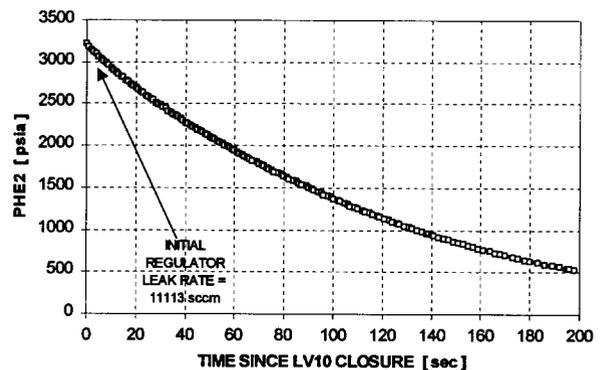


Figure 32. PHE2 Decay Curve at Pre-DSM Press.

One additional regulator leakage assessment was possible, after the DSM maneuver itself. Within uncertainties, it demonstrated the same leak rate as the pre-DSM pressurization. Figure 33 is the post-DSM analogue to Figure 31, showing the

bipropellant tank filling during the leaking regime of the regulator, up until the closure of LV10. For consistency with Figure 31, a second-order curve fit was used; however, the behavior is essentially first-order, as can be verified in Figure 33.

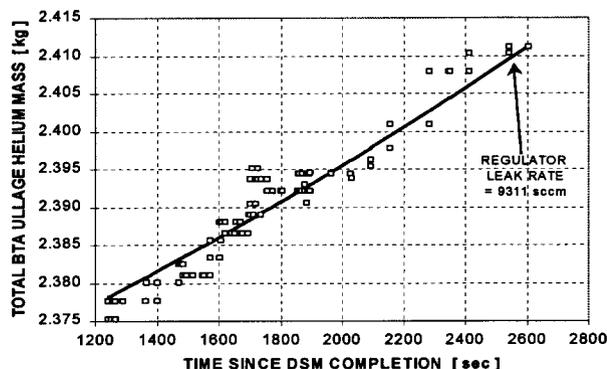


Figure 33. BTA He Mass During Post-DSM Press.

The apparent regulator leak rate from Figure 33 is 9311 sccm (558660 scch), but actually represents a larger leak than the pre-DSM pressurization, because the upstream pressure decreased from 3200 psia to 2300 psia across DSM. However, it is quite possible that the regulator leak rate (or, equivalently, particle size) did not change across DSM, within uncertainties.

Following DSM, the regulator outlet pressure showed an anomalously high value. This was explained as a consequence of closing LV20 and LV30 ten and twenty seconds, respectively, after the closure of LV10. For regulated, main-engine burns longer than 60 minutes, it is desirable to close LV20 and LV30 as soon as possible, to avoid possible cold traps for NTO migration and condensation. Both LV20 and LV30 have forward pressure-relief features, and this forward-relief feature was inadvertently tested following DSM. In the presence of a leaking regulator, it would be theoretically possible to overpressurize the regulator outlet if LV20 and LV30 were closed too early and did not forward relieve.

Figure 34 is the post-DSM analogue to Figure 32. Note that the regulator outlet pressure is included as well. An independent anomaly in data playback led to poor time resolution in the data of Figure 34. This exacerbated attempts to determine if the regulator outlet was overpressurized following LV30 closure. Note again that the PHE2 decay curve follows an exponential trend, as expected. This

exponential curve fit was used to find the helium flow rate at the time of LV30 closure. Fortunately, the first regulator outlet pressure telemetry sample was only eight seconds after LV30 closure, and it showed a pressure value consistent with the forward relief of LV20 and/or LV30. Even in the pathological case of having LV20 and/or LV30 relieve precisely at this point (eight seconds after LV30 closure), as required by regulator pressure telemetry, the maximum regulator outlet pressure would have only been 657 psia, slightly below the proof pressure of 685 psia. There is no reasonable expectation that LV20 or LV30 failed to pressure relieve as designed, but this worst-case assessment allays any fears that the regulator might have been damaged post-DSM.

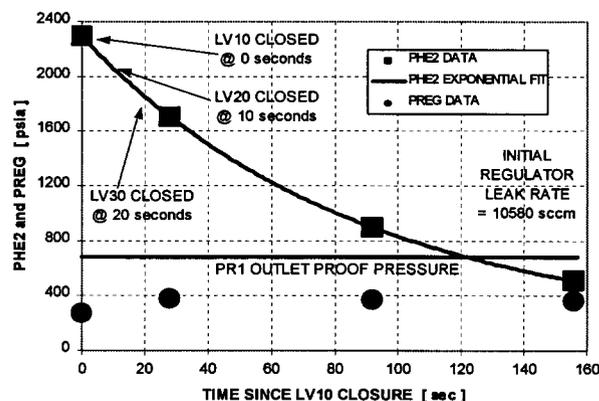


Figure 34. PHE2 Decay Curve Following DSM

LV10 was cycled open and closed twice more during the Cassini mission to date, between TCM-9 and TCM-10 and during TCM-13. However, it was impossible to determine the regulator leak rate during these activities, because LV10 was closed while the regulator was still in its filling regime (rather than its leakage regime, above the expected regulator lock-up point). Regulator function will be tested five weeks before SOI with a large, deterministic main-engine maneuver, TCM-20. This TCM will target Saturn's interesting moon, Phoebe, and it offers excellent risk mitigation for the Cassini PMS before its required use during the critical SOI burn.

## X. Conclusions

The Cassini Propulsion Module Subsystem (PMS) has performed excellently throughout over four years of mission operations. Consumable usage has typically been better than expected, and there are

no PMS consumable limit violations expected, even for an extended mission.

TCM performance has been excellent, and the propellant margin continues to increase over time. There is no evidence for pressure transducer drift, in contrast to the Galileo, Voyager, and TOPEX-Poseidon missions. Cassini has allowed a thorough characterization of 0.9-N thruster pulse-mode performance over a wide range of thruster on-times. Cassini helium budgets have demonstrated no internal or external helium leakage in both the monopropellant and bipropellant portions of the PMS.

The only PMS anomaly of note during the Cassini in-flight mission to date is the presence of a very large primary regulator leak. This leakage is most likely due to particulate contamination. Surprisingly, it has little consequence for the Cassini mission, due to the presence of a high-pressure latch valve that can be used to control the timing of bipropellant tank pressurization. This is despite the fact that the regulator leak is three to four orders of magnitude higher than the spec leak rate.

The prospect for Cassini completing its prime mission at Saturn are excellent, particularly since almost all engineering activities required for orbit insertion and probe relay have been demonstrated in flight. There will likely be resources for an extended mission as well, allowing for further characterization of the ringed planet, its rings, intriguing moons, and magnetosphere.

## **XI. Acknowledgments**

The research described in this paper was carried out by Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. The authors wish to acknowledge the following individuals who contributed to this research: Carl Guernsey, Kevin Johnson, Jeremy Jones, Kerylyn Lay, Allan Lee, Michael Leeds, Hartwell Long, Earl Maize, Barry Nakazono, Tim O'Donnell, Morgan Parker, Laura Sakamoto, David Skulsky, Erica Spurgeon, Julie Webster, and Jeff Weiss. The authors would like to gratefully acknowledge the contributions of Marisol Carranza, Arbi Karapetian, and Genevieve Velarde, who helped bring this paper to fruition.

## **XII. References**

- <sup>1</sup> Barber, Todd J.; Krug, F. A.; and Froidevaux, B. M.: "Initial Galileo Propulsion System In-Flight Characterization," AIAA Paper 93-2117, June, 1993.
- <sup>2</sup> Barber, Todd J.; Krug, F. A.; and Renner, K. P.: "Final Galileo Propulsion System In-Flight Characterization," AIAA Paper 97-2946, July, 1997.
- <sup>3</sup> Leeds, Michael W.; Eberhardt, R. N.; and Berry, Robert L.: "Development of the Cassini Spacecraft Propulsion System," AIAA Paper 96-2864, July, 1996.

