

# Mission Summary of Cassini Spacecraft Guidance and Control Hardware Health and Performance

Joan Stupik\*

*Jet Propulsion Laboratory, California Institute of Technology*

The Cassini-Huygens mission ended on September 15, 2017, after nearly two decades in flight. The well-designed Cassini spacecraft had robust hardware that permitted two extended missions, lasting nine years longer than the expected prime mission. At the end of the mission, the Attitude and Articulation Control Subsystem (AACS) was using two pieces of redundant back-up of hardware, one reaction wheel and the hydrazine thruster branch, due to hardware anomalies earlier in the mission. The back-up hardware performed nominally through the rest of mission. The prime reaction wheels at the end of the mission had reached more than 130% of the consumable limit for number of revolutions. No thruster on either thruster branch accumulated more than 45% of the consumable limits. The inertial reference unit slightly exceeded the pre-launch requirements on bias error, but as the software continuously estimated this value in flight, the attitude estimation was not adversely affected. The star trackers performed nominally, and though there was a spacecraft anomaly in 1998 related to the star trackers, the origin was not in hardware itself. The Sun sensors and accelerometer both performed as expected and met all requirements throughout the mission. Ultimately, the lifetime of the Cassini spacecraft was not limited by hardware performance. Planetary protection requirements necessitated the end of the mission as the spacecraft's propellant reserves depleted, and Cassini plunged into Saturn's atmosphere with a healthy attitude control system.

## Nomenclature

*AACS* Attitude and Articulation Control Subsystem  
*ACC* Accelerometer  
*HGA* High Gain Antenna  
*OTM* Orbit Trim Maneuver  
*IRU* Inertial Reference Unit  
*RCS* Reaction Control System  
*RWA* Reaction Wheel Assembly  
*SET* Single Event Transient  
*SOI* Saturn Orbit Insertion  
*SRU* Stellar Reference Unit  
*SSA* Sun Sensor Assembly  
*TCM* Trajectory Correction Maneuver

## I. Introduction

Cassini was one of the largest and most complex interplanetary spacecraft that has been constructed and launched. The spacecraft was launched in 1997. The Cruise phase of the mission was seven years long, during which time Cassini made gravity assist flybys of Venus twice, Earth, and Jupiter on the way to Saturn. The Tour phase of the mission began when the spacecraft successfully entered into orbit around Saturn in 2004. Cassini completed its prime mission (between 2004 and 2008) and was subsequently approved for first

---

\*Guidance and Control Engineer, Guidance and Control Section, Mail Stop 198-235, 4800 Oak Grove Drive, Pasadena, CA 91109-8099.

and second extended missions. Cassini was a three-axis stabilized spacecraft that carried a payload of 12 instruments, including antennas, cameras, and spectrometers. The bulk of these instruments were fixed with respect to the spacecraft body and lacked the ability to articulate.<sup>1</sup> Therefore, the Attitude and Articulation Control Subsystem (AACS) slewed the whole spacecraft to point an instrument to a science-specified attitude.

The Cassini mission concluded on September 15, 2017, after the bi-propellant reserves were nearly depleted. To prevent the possibility of undesirable contamination of the moons Titan and Enceladus, Cassini was incinerated in Saturn's atmosphere. Before entering the Saturnian atmosphere, Cassini ventured into previously unexplored territory - diving through a 1900 km (1200 mile) gap between the inner edge of the rings and the top of Saturn's atmosphere. A total of 22 dives were executed before entering Saturn's atmosphere on September 15, 2017. These concluding orbits provided the opportunity to make unique and unprecedented observations. The relatively aggressive finale of this NASA flagship mission was made possible by the continued nominal and near-nominal performance of spacecraft hardware. This paper summarizes the health of Cassini's Attitude and Articulation Control Subsystem (AACS) hardware at the end of the mission, and discusses the performance of this hardware over two decades of flight. Also included in the paper are several topics unique to each piece of hardware, regarding calibrations, spacecraft anomalies, and ground operation procedures.

Cassini's AACS sub-system used star trackers (SRUs) and gyroscopes (IRUs) to determine the spacecraft's celestial attitude. For the majority of the tour, attitude control was done using reaction wheels, to meet the fine science pointing requirements. Small hydrazine thrusters were used for a number of events during tour, including momentum biasing,  $\Delta V$  maneuvers less than 0.3 m/s, and to maintain control authority when flying through the atmosphere of both Titan and Saturn. The AACS team was responsible for the hydrazine thrusters' valve drive electronics. Cassini also had digital Sun sensors assemblies (SSA) to provide a Sun direction estimate when the Sun was in the field of view of the sensor. The SSA data was not nominally used in the attitude estimator - the Sun sensors measurements were used in fault monitors, and would be used for attitude reinitialization in the event that the spacecraft lost inertial attitude knowledge. Cassini also had one accelerometer that was used only for  $\Delta V$  maneuvers executed with the main engine. Other AACS hardware included engine gimbal actuators, which controlled the pointing of the main engine thrust-axis, and the AACS flight computer, which was distinct from the main spacecraft computer. Most of the AACS sensors and actuators had both prime and redundant units.

Fig. 1 is a diagram of the Cassini spacecraft. The high gain antenna (HGA) was a parabolic reflector dish that was four meters in diameter, and was mounted along the spacecraft body -Z-axis. It was used to communicate with Earth. The IRUs, SRUs, sun sensors, one of the reaction wheels, and one of the thruster clusters are labeled.

In this paper, the reaction wheel drag torque trends are shown over the course of Cassini's tour (reaction wheels were used minimally during the 6.7-year Cruise phase). Descriptions are also given regarding the ground tools used to manage the on-board momentum and the cage instability that appeared in one wheel bearing early in the mission, which drove a swap for the backup wheel. For the hydrazine thrusters, trending of the individual thrust is provided since 2007. Visible in these trends is the degradation that occurred on two of the thrusters in the prime thruster branch, driving a change to the back-up thruster branch in 2009. The IRUs performed nominally for the entirety of the mission, except for occasional spikes in the rates due to impacts by high-energy particles. Information from the IRU calibrations is given, showing the misalignments and scale factor errors over time. The star trackers, or stellar reference units (SRUs), also performed well during the mission. A spacecraft anomaly occurred when a fault protection monitor related to the SRUs triggered. However, this was not a fault in the SRU hardware itself, but in a fault protection monitor threshold being set too tight. The SRU did require some careful management of bright bodies in the FOV, a summary of which is given in this paper. The Sun sensors had little trending data available (the available data is given in this paper) as periodic calibrations were not scheduled for these sensors, but the sensors were checked after each time that they were exposed in a region where they might be susceptible to micro-meteoroid hits. The accelerometer bias was calibrated before each main engine maneuver that Cassini performed, and the drift of the bias throughout the mission is trended. The main engine gimbals and the AACS flight computers are not described in this paper. The performance of both pieces of hardware was nominal. The main engine gimbal nominal performance is evident in the excellent maneuver execution performance, given in Ref. 3.

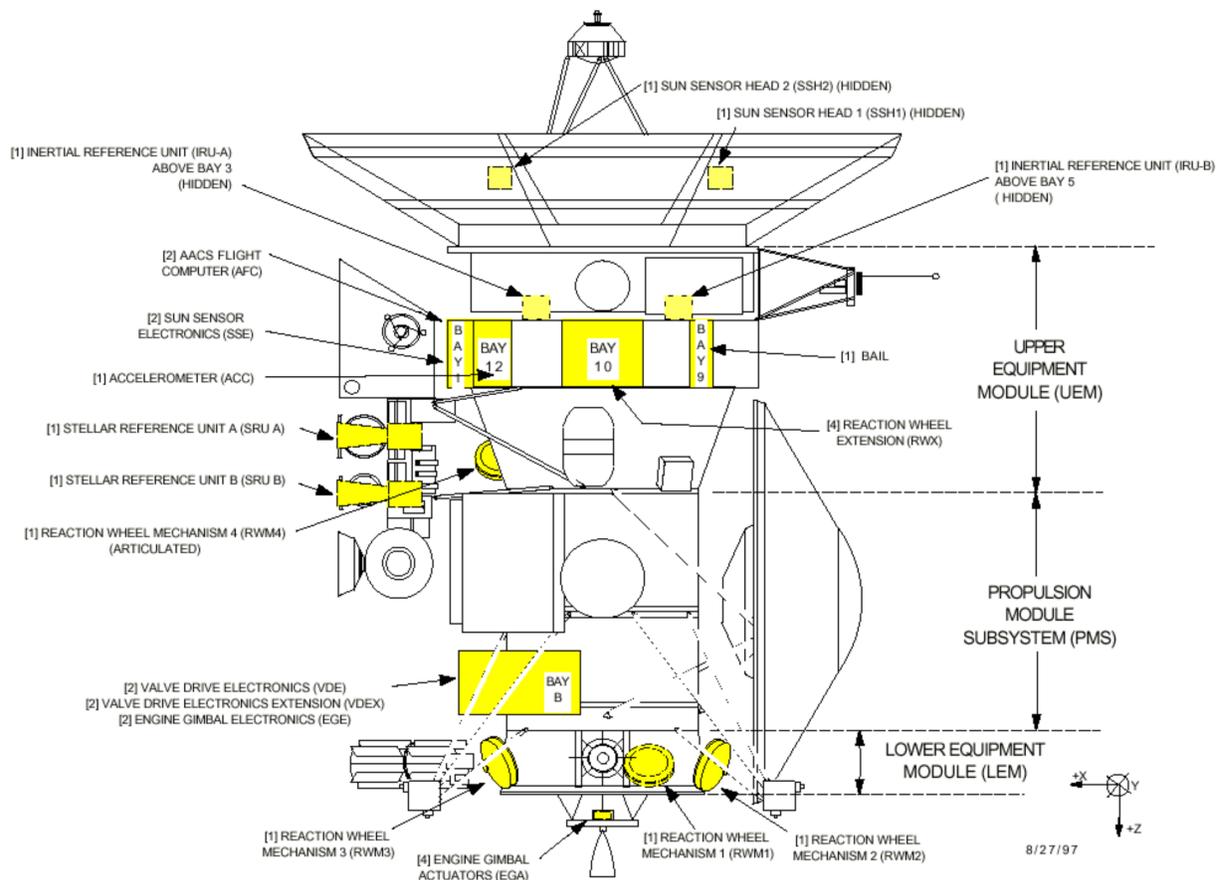


Figure 1. Cassini Spacecraft Diagram<sup>2</sup>

## II. Reaction Wheels

The Cassini spacecraft was three-axis stabilized, using either reaction wheels (RWAs) or thrusters for attitude control. The majority of the science instruments were fixed to the spacecraft body and therefore required the whole spacecraft to slew in order to point to an attitude specified by a science team.<sup>1</sup> Over the course of the mission, Cassini's reaction wheels performed well. Science pointing requirements<sup>2</sup> were met, and Cassini dove into Saturn's atmosphere with three healthy reaction wheels.

Reaction wheels did not directly consume propellant as the thrusters did, and the pointing control and stability achievable in the RWA control mode was approximately 50 times more accurate than the capability of the thrusters. The thrusters were used for several specific events (discussed in Section III), but the vast majority of Cassini's tour at Saturn was completed in the RWA control mode. To preserve the health of the reaction wheels for the tour, Cassini flew in the thruster control mode for the majority of Cruise, leaving the reaction wheels powered off.

Cassini carried a total of four reaction wheels, but only used three at once for active control. The fourth wheel provided redundancy. Each reaction wheel could store a maximum of approximately 36 Nms and run with a maximum wheel speed of about 2020 rpm. The configuration and mounting of the RWAs on the Cassini spacecraft is shown in Fig. 2 below.

As can be seen in Fig. 2, three of the reaction wheels are mounted orthogonally. RWA4 (the fourth reaction wheel) was mounted on a platform that could be articulated using a motor. As RWA4 could be articulated<sup>4</sup> to match the alignment of any of the three other reaction wheels, Cassini flew with a fully redundant spare reaction wheel. The spacecraft launched with RWAs 1, 2, and 3 as the prime set. During the limited time spent on reaction wheel control during Cruise, RWA3 began exhibiting signs of cage instability.<sup>2,5</sup> In 2003, RWA4 was articulated to match the orientation of RWA3, and replaced it in the set of active reaction wheels. Cassini flew the entirety of tour with a prime reaction wheel set of RWAs 1, 2, and 4.

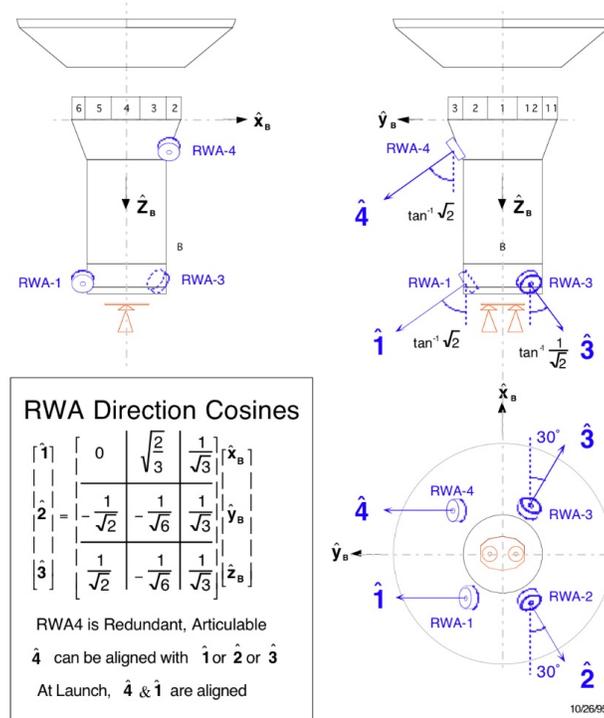


Figure 2. Reaction Wheel Locations and Orientations<sup>10</sup>

The reaction wheels, though not directly consuming propellant, did frequently require the use of the thrusters to manage RWA momentum. Cassini flew with intentional, non-zero momentum biases. The goal in flying with a non-zero momentum bias in the wheels was to minimize the amount of time that the reaction wheels spent spinning below a specific low rate (in the “low-rpm” region). The thickness of the elasto-hydrodynamic lubrication film between the bearing balls and the bearing races was related to the spin rate. At low rpm (on Cassini, less than 300 rpm), the thickness could have been so small that there was metal-to-metal contact between the balls and the races. This contact promoted wear and heating of the lubricant, leading to polymerization of the lubrication. During Jupiter flyby, fault protection detected excessive friction torque in RWA2. This was due to prolonged dwell near zero-rpm. The automated response was to power off RWA2 and switch to RCS control.

The momentum bias stored in the reaction wheels was managed with frequent momentum biases using the reaction control thrusters. The biases were placed to minimize the time in the low-rpm region while supporting the required science slews and tracking. The method of managing the momentum, the “Reaction Wheel Bias Optimization Tool”, is outlined in Section II. B.

In flight, it was discovered that the reaction wheels were, on occasion, adversely affecting some science observations through microphonics. The vibrations of the reaction wheels at specific spin rates were transferred through the spacecraft structure to the Composite Infrared Spectrometer (CIRS), and disrupted the instrument’s internal alignment enough to occasionally affect the observations. The in-flight mitigation strategy is discussed in Section II. C.

## A. Hardware Performance

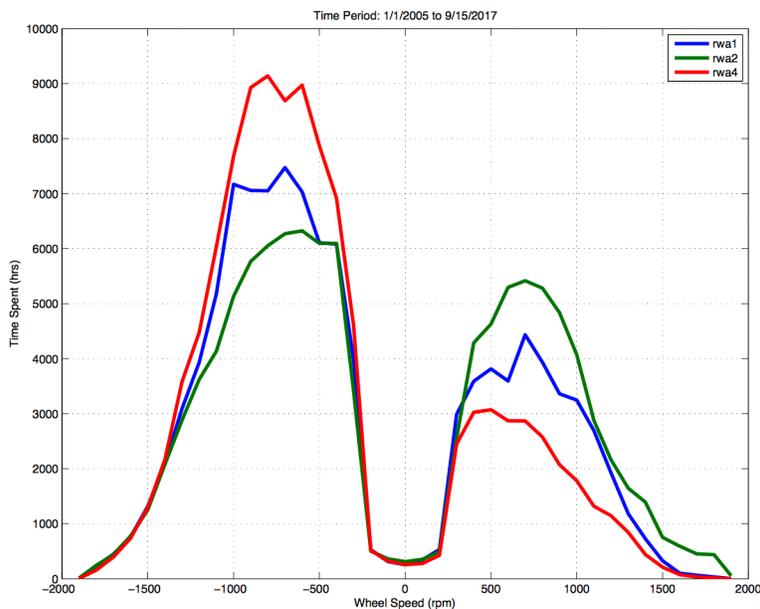
Cassini had a set of functioning reaction wheels that met the science pointing requirements available for the entirety of the mission. RWA3 was retired as a prime wheel during Cruise, but RWA4 performed nominally for the rest of the mission. As part of operations, the Cassini AACS team tracked the RWA consumables of: number of revolutions, time spent in low-rpm, and on-off cycles. The consumable limits were guidelines, but it was not guaranteed that the wheels would perform nominally when below the limits (as evidenced by RWA3). The drag torque was a good indicator of RWA health and was also closely monitored during the

mission. The final consumables are given in Table 1.

**Table 1. Reaction Wheel Consumables**

Consumable	Flight Allocation (per wheel)	RWA1	RWA2	RWA3	RWA4
Revolutions (billion revs)	4	146%	148%	14%	133%
Low-rpm Time (hours)	12000	70%	69%	28%	40%
On-Off Cycles	450	76%	87%	28%	85%

As the mission had two extended missions spanning nine years, it is not surprising that the revolutions budget was exceeded for all wheels except for RWA3 (which spent most of the mission powered off). Even with the extended missions, however, the low-rpm time and the on-off cycles remained within the allocation. The low-rpm time consumable was carefully managed by the operations team using a ground software tool optimization process (see Section II. B). For Cassini, the low-rpm region was defined as less than  $\pm 300$  rpm. The efforts to minimize the time spent in this region are clearly visible in the history of the reaction wheel spin rates since Saturn Orbit Insertion (SOI) in Fig. 3, where there is a notable drop-off at  $\pm 300$  rpm. Very high spin rates were also avoided, to slow the accumulation of revolutions. The result in Fig. 3 shows two lobes in the spin-regions where the reaction wheels spent the most time, between 300 rpm and 1500 rpm both in the positive and negative directions.



**Figure 3. Reaction Wheel Time Spent at Spin-Rates**

In general, the reaction wheels spent more time spinning in the “negative” direction. This was due to the direction of frequent downlink rolls (during communication with Earth, the spacecraft would roll about the high-gain antenna’s axis, to make measurements with the fields and particles science instruments). The rolls were always in the same direction, resulting in the bias toward the negative spin rates seen in Fig. 3.

Fig. 4 shows the history of the drag torque in each of the three wheels that have been in use since SOI. The solid black line in the RWA1 plot represents the predicted drag torque. RWA2 and RWA4 remained close enough to the predicted drag torque that the line is obscured in those plots. The drag torque shown is the average drag torque across three months in discrete RWA spin rate bins. The cooler colors represent the older telemetry, while the warmer colors are more recent. RWA2 and RWA4 both show lobes of higher drag torque at low spin rates (specifically in the  $\pm 300$  rpm, low-rpm region) for a period of a few years, but not in the last years of the mission. These “lobes” are a result of spikes in the drag torque when the reaction wheels were spinning at lower rates. RWA2’s lobe spans from 2008 to 2010, while RWA4 spans from 2009 to

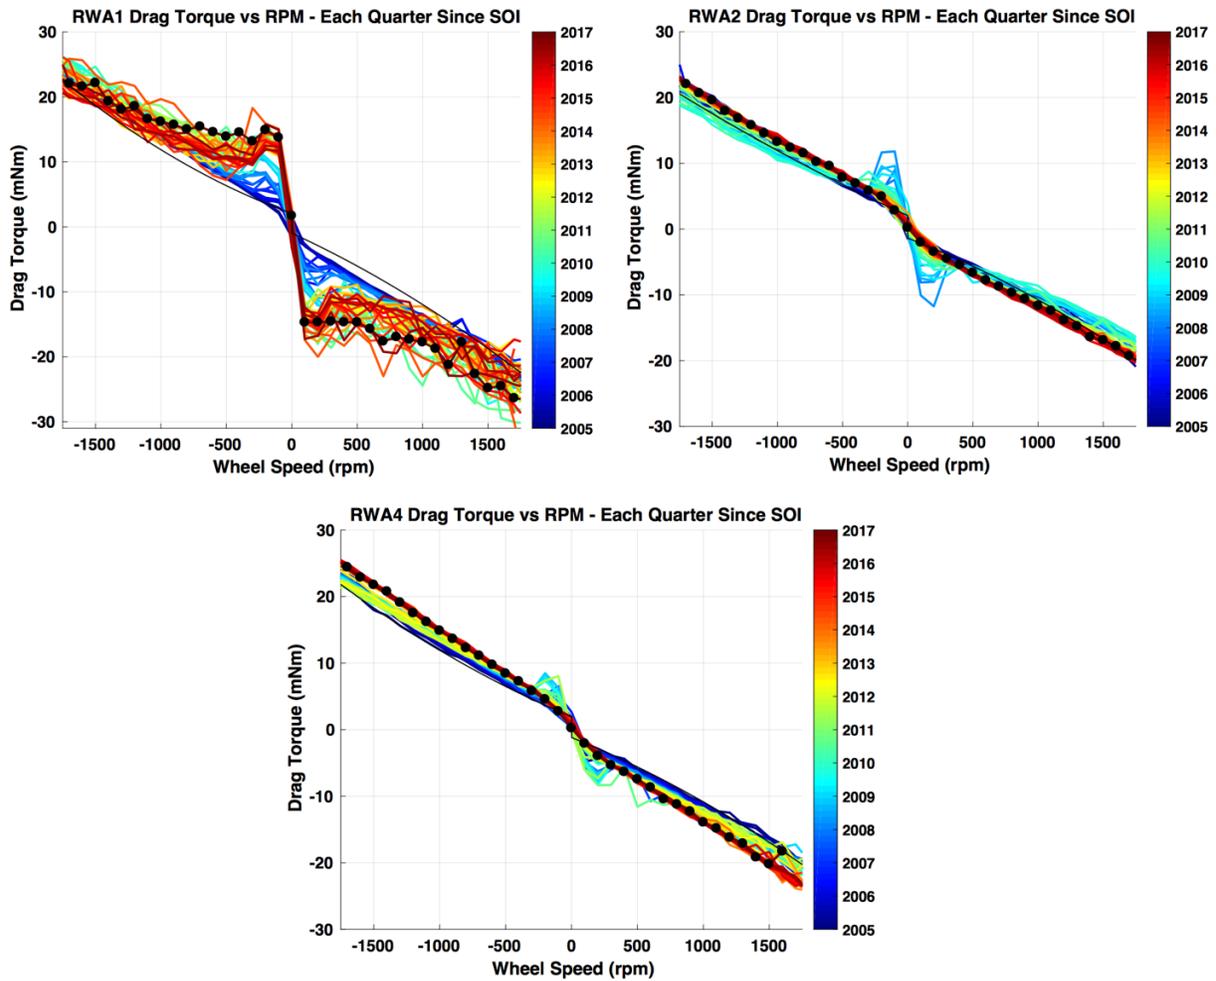


Figure 4. Reaction Wheel Drag Torque History

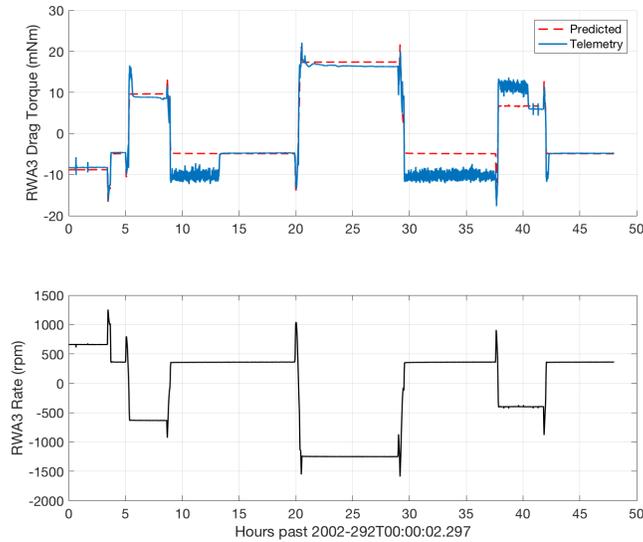
2011. However, since 2011, the drag torque trends of both RWA2 and RWA4 have been well behaved, with neither a notable increase since 2012 or 2013, nor evidence of drag torque spikes at lower spin rates.

RWA1 consistently had drag torque that was higher than the prediction - exhibiting the drag torque spike behavior when the wheel was spinning in the low-rpm region - all the way through the end of the mission. These torque spikes were sometimes visible in the spacecraft pointing, but did not negatively impact science pointing.

Despite elevated drag torque in RWA1, the prime set of wheels for most of the mission (RWA's 1, 2, and 4) ensured that science pointing requirements were met during tour. A problem did arise with RWA3 during Cruise, prompting the swap out of RWA3 for RWA4.

RWA3 began to exhibit signatures associated with cage instability during Cruise, in 2002.<sup>6</sup> This uncontrolled vibration of the bearing cage manifests in telemetry as elevated, noisy drag torque, as seen in Fig. 5. This figure shows the first observances of the RWA3 cage instability, beginning at approximately 2002-DOY292T09. The top plot of Fig. 5 shows both the predicted and in-flight drag torques, while the bottom plot shows the spin rate of RWA3. After a turn is completed at 2002-DOY292T09, RWA3 is at a constant spin rate until 2002-DOY292T20. However, after turn completion, the drag torque in the top plot shows sudden, noisy elevation above the prediction for a few hours, before dropping back to the predicted level. The same signature was seen again the following day (beginning around 2002-DOY293T06 in the figure). The drag torque oscillation persisted through a turn at 2002-DOY293T14 before falling back to the prediction at 2002-DOY293T17.

After RWA3 continued to exhibit this behavior, the decision was made in 2003 to replace RWA3 with



**Figure 5. First Observance of RWA3 Bearing Cage Instability**

RWA4 in the prime set. Also, project engineers significantly curtailed the use of the RWAs to preserve them for use at Saturn. The wheels provided fine pointing control and stability that was required for observations at Saturn to meet the mission science objectives. The spacecraft continued to use the prime set of reaction wheels 1, 2, and 4 for the rest of the mission, with the exception of a few hours in November of 2010 and the months of March and April in 2011. In November 2010, RWA3 was reintroduced as a prime wheel for a period of 50 hours to look for similar cage instability symptoms. None were observed during the 2010 test, and RWA3 was again reintroduced as a prime wheel for the science observation sequence spanning March and April of 2011. In late March, the drag torque oscillation signatures reappeared, and RWA4 again replaced RWA3. RWA3 remained the back-up wheel for the remainder of the mission. For a more detailed discussion of RWA3, see Ref. 6.

## B. Reaction Wheel Biasing Optimization Tool

To help maintain RWA health, more frequent biasing of RWA momentum was instituted. Mission planners reserved 90-minute RWA momentum bias candidate windows immediately before each nine-hour Deep Space Network track. Cassini was always quiescent and Earth-pointed during these windows. A new ground software tool, the Reaction Wheel Bias Optimization Tool (RBOT),<sup>5</sup> was developed. The AACS team used RBOT to decide which windows should be utilized for a RWA momentum bias, and what the initial momentum state should be.

Given a selection of opportunities to insert a momentum bias, the ground software tool identified the optimal location and size of momentum biases. The optimization rejected bias strategy solutions that cause the reaction wheel rates to exceed the flight limit. Each of Cassini's reaction wheels had a total momentum storage of 36 Nms, with a maximum spin rate of approximately 2020 rpm.<sup>6</sup> In flight, the spin rates were kept under 1850 rpm to protect against momentum build-up due to unmodeled external torque. The cost function for the optimization penalizes wheel rates under the low-rpm limit (to preserve the wheel bearing health) and lightly penalizes very high spin rates (to limit the number of total revolutions).<sup>5</sup> The overall effect was that the reaction wheels spent most of their time at rates between 300 and 1500 rpm (as can be seen in Fig. 3).

The RBOT tool preserved the health of the wheels while keeping the operations work manageable. Science observation periods were organized into sequences, typically seven to ten weeks long. Each sequence would undergo development and testing as a unit. The science observations in each sequence were planned many months in advance of execution. After the majority of the science observations in a sequence were designed, the engineering team would add engineering activities and test the complete set of commands. Well in

advance of the science observation planning, the project would set aside times for potential momentum bias opportunities. Once the science observations were planned and the spacecraft attitude profile for those weeks was determined, the RBOT tool selected from the reserved bias opportunities and calculated the optimal initial momentum state to minimize the time spent in low-rpm for the period before the next bias. A later version of RBOT included an additional feature to minimize propellant use along with low-rpm time, by choosing a smaller momentum bias when possible. For more information about RBOT, see Ref. 5.

### C. Reaction Wheel Microphonics Impact on CIRS

After launch, the Composite Infrared Spectrometer (CIRS) observed degraded data during some of their observations. The instrument relied on ultra-precise alignment between the internal mirrors and frequency generator. Before launch, there was expected to be sufficient damping to prevent reaction wheel vibrations from disrupting CIRS - unfortunately, despite the damping, CIRS would lose data in flight when the RWAs spun at certain rates. These “keep-out” zones were mostly above 1300 rpm and applied only during CIRS observations. Operationally, it was impossible to ensure that all the reaction wheels were spinning at acceptable rates during every CIRS observations. To address the issue of RWA microphonics, a process was developed between the CIRS team and the AACS team whereby the CIRS team could request mitigation for certain pre-identified high-priority CIRS observations.

For a spacecraft with three active reaction wheels, the only ways to change the wheel rates at any given time are (1.) change the spacecraft momentum state or (2.) change the spacecraft attitude. Changing the momentum state is theoretically achievable by managing the momentum bias strategy, either by adding more biases (consuming fuel and introducing more  $\Delta V$ ) or changing the momentum state of an existing bias (incurring more low-rpm time or revolutions, as the RBOT optimal solution is abandoned). Changing the spacecraft attitude could be done by rotating the spacecraft about the instrument boresight (the CIRS instrument was a slit, and therefore potentially sensitive to these rotations). Fig 6 shows the flow of the CIRS RWA Mitigation Process.

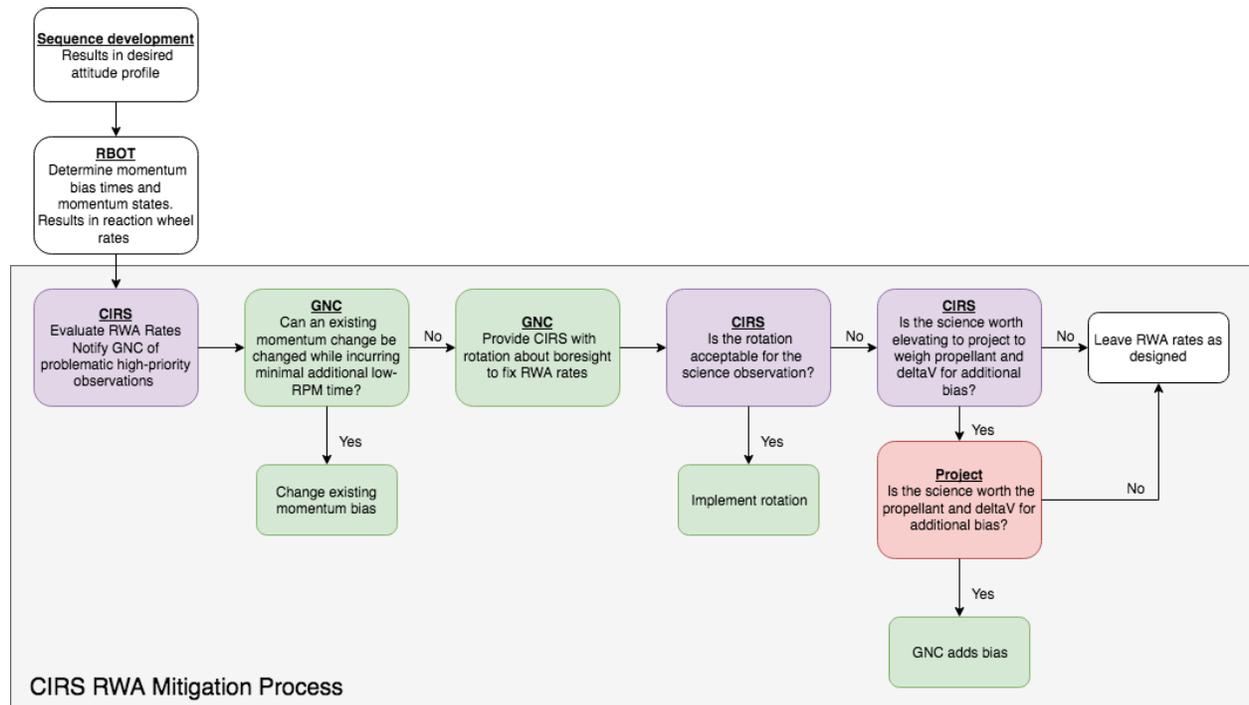


Figure 6. CIRS Reaction Wheel Mitigation Process

For operational feasibility, the CIRS team was only permitted to invoke the mitigation process for less than 15 high-priority observations per science sequence, identified in advance of the reaction wheel momentum bias placement process. After the AACS team developed the initial momentum bias strategy, the CIRS team evaluated the predicted reaction wheel rates and identified any of the high-priority observations for which

they wanted to invoke the mitigation process. The AACS team’s order of priority for mitigation strategies was first to attempt to find a slightly less RBOT-optimal wheel rate profile by changing the momentum state at the end of the previous bias, second to search for a rotation about the CIRS boresight, and lastly to add a new momentum bias (which required the approval of project management).

This process was successful in mitigating the unexpected interaction between the RWAs and the CIRS instrument without a significant impact to operations cost.

### III. Reaction Control Thrusters

The reaction control system (RCS) consisted of two independent branches of eight hydrazine thrusters each (total of sixteen thrusters). The branches were called A-branch and B-branch. On each branch, there were four thrusters oriented along the spacecraft body Z-axis (providing control in the X- and Y-axes) and four along the Y-axis (providing control in the Z-axis), as shown in Fig. 7. The arrows indicate the directions of the reaction forces imparted on the spacecraft by the firing thrusters. At launch, the thruster force was approximately 1.0 N. By the end of the mission, the decreasing tank pressure caused the thrust to reach a minimum of 0.57 N.

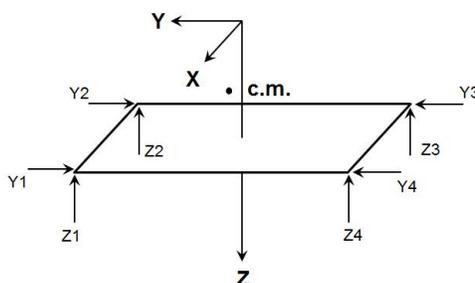


Figure 7. Representative RCS Thruster Branch

While the spacecraft spent the majority of the tour on reaction wheel control, the RCS thrusters were used to control the spacecraft attitude during Cruise, to preserve the life of the reaction wheels for the exploration of the Saturnian system. The thrusters also had several uses during the tour at Saturn. First, they were used for small (less than 0.3 m/s)  $\Delta V$  burns, called trajectory correction maneuvers (TCM) during the Cruise phase of the mission and orbit trim maneuvers (OTM) during the tour. Cassini had a 450-N main engine that was used for burns larger than 0.3 m/s (there was a back-up main engine which was never fired). Second, the thrusters were used to bias the momentum stored in the reaction wheels. Momentum biases occurred approximately five times per orbit.<sup>7</sup> Finally, the spacecraft transitioned to the RCS control mode when the science objectives required the greater torque authority provided by the thrusters. These times included flying through Titan’s atmosphere, the plumes of Enceladus, Saturn’s atmosphere during the Grand Finale, and also times when science observations required fast turn capability.

One of Cassini’s thruster branch, A-branch, had some unexpected degradation in two of its Z-facing thrusters in 2009. After a significant under-burn on a RCS OTM, the problem was identified and the whole of A-branch was swapped out for B-branch. The possible causes of the degradation are explored in Ref. 8. After the swap, the health of the B-branch thrusters was carefully monitored and managed. A new trending tool was developed, and a new method of momentum biases was utilized to more evenly distribute the pulse and throughput consumables across all thrusters in the branch. Cassini ended the mission with B-branch as the prime thruster set, and saw no signs of similar degradation despite comparable usage levels.

#### A. Hardware Performance

The mission began with A-branch as the prime (or active) branch. In 2009, anomalous degradation in thrusters Z3 and Z4 necessitated a switch to B-branch as the active branch. The total thruster cycles (or pulses) and throughput are listed in Table 2 for both branches. A-branch was prime for 12 years (1997–2009) while B-branch was prime for the remaining 8 years of the mission. It is therefore not surprising that, over the course of the mission, A-branch accumulated more thruster cycles and propellant throughput than

B-branch. This is seen in Table 2.

**Table 2. Average Thruster Branch Consumables**

Consumable	Flight Allocation (per thruster)	Branch	Average	
			Y	Z
Throughput (kg)	25	A	16%	36%
		B	16%	29%
Thruster Cycles	273,000	A	18%	41%
		B	21%	25%

Table 2 also shows that the B-branch thrusters accumulated the cycles and throughput more evenly across the Y- and Z-thrusters when compared with A-branch. On average, the A-branch Z-thrusters accumulated 20% more of the throughput consumable limit than the Y-thrusters, and 23% more of the thruster cycles consumable limit. These differences are smaller on B-branch – the Z-thrusters accumulated an average of 13% more of the throughput consumable limit than the Y-thrusters, and an average of 4% more of the thruster cycle limit. This reduction in disparity between Y- and Z- thruster usage was purposefully accomplished through operations. To preserve the health of B-branch, a new type of momentum biasing was designed to isolate thruster firing only to the Y-thrusters during the momentum change. This bias type was called “y-biasing” and is discussed in Section III. B.

Due to a lack of visibility, the degradation in the performance of A-branch was not observed until an OTM did not achieve the desired change in velocity, resulting in a  $\Delta V$  penalty of 5.7 m/s spread over subsequent maneuvers.<sup>8</sup> In response to the unobserved degradation in the thrust, the new bias strategy (y-biasing) was utilized and a thruster performance trending ground tool was developed. The estimation tool calculated the thrust of each individual thruster from related spacecraft telemetry.<sup>9</sup> For more information on the estimation tool, see Ref. 9. The mission thruster performance of the Z-thrusters, both A and B branches, is shown in Fig. 8. Similarly, performance of the Y-thrusters is shown for both branches in Fig. 9. Note that some degradation in the thrust is expected as the pressure in the tank decreases over the course of the mission. The red dotted lines in both of these figures represent the expected thrust, based on the pressure in the tank at that time.

The solid black line in Fig. 8 represents the date of the OTM that failed to achieve the desired  $\Delta V$  and alerted the operations team to the dramatic degradation on Branch-A, Z3 (AZ3) and less dramatic degradation on Branch-A, Z4 (AZ4). The A-branch trending plot was created with the new trending tool after the swap to B-branch, to demonstrate that the degradation was visible in the reconstructed data. The swap to B-branch did not occur until after the degradation had been reviewed and the action plan approved, which was several months after the OTM represented by the black line. The degradation continues drastically on AZ3 until the swap. By the time of the swap in early spring of 2009, the degradation of AZ4 is more apparent. The B-branch thruster trending in Fig. 8 shows that no similar signs of degradation were observed on B-branch thrusters. B-branch remained the prime thruster branch for the remainder of the mission and performed nominally throughout.

The Y-thrusters on both thruster branches performed nominally for the entire mission, as seen in Fig. 9 when trending was available. Prior to 2007, the telemetry channel containing the on-time of the thrusters was recorded at a very low frequency. The on-time telemetry is one of the crucial inputs to the thrust reconstruction tool<sup>9</sup> - the limited availability of this telemetry channel for the first several years of the mission means that the trending information available for the thrusters is only available since 2007. The A-branch Y-thruster thrust levels appear in Fig. 9 to fall just below the expected thrust envelope, while still following the expected trend (the cause of the offset is thought to be a result of a mis-representation of the thruster rise and/or fall time<sup>9</sup>). As thrust degradation would be most visible as a deviation from the trend, the information available for the A-branch Y-thrusters is sufficient to show that there was no sign of the degradation seen by the Z-facing thrusters.

## B. Y-Biasing

During an RCS OTM (an OTM using the RCS thrusters instead of the main engine), the Y-thrusters control the roll about the burn axis. The Z-thrusters provide the thrust required by the navigation team to maintain

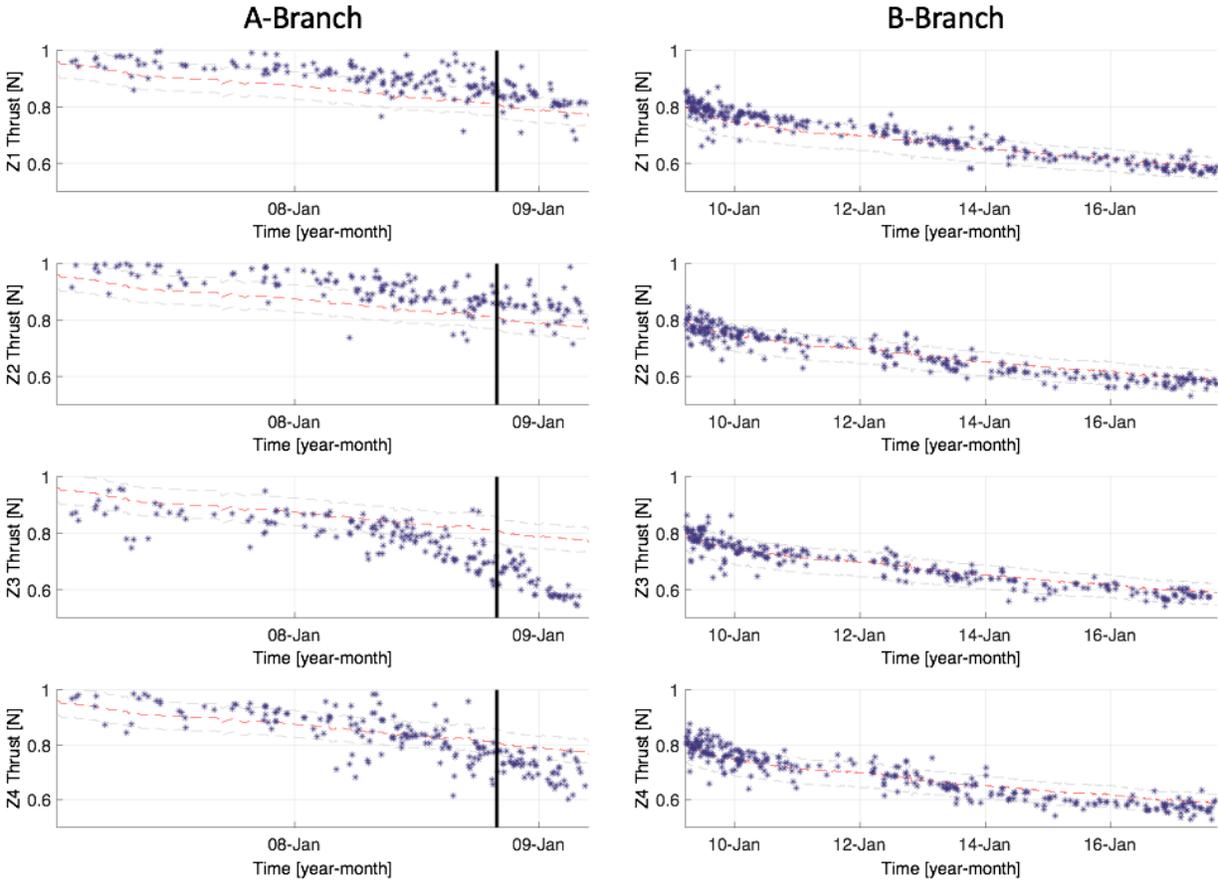


Figure 8. RCS Z-Thruster Branches Mission Performance.<sup>9</sup> A-Branch (left) and B-Branch (right) Z-Thrusters Thrust Magnitude Estimates. Minimum on-time of 5 seconds. Solid (red) line shows expected thrust. Dash-dot (black) lines show  $\pm 5\%$ , an envelope of expected deviation. Asterisk (blue) points are thrust estimates. Solid (black) vertical line indicates the time of the OTM in October of 2008 that alerted the team to the thrust degradation.

the mission trajectory, and therefore see more usage. The distribution of thruster use during fast science turns and maintaining control authority in the presence of external torque from atmosphere was dependent on attitude - which was dictated by science needs. The best place to reduce Z-thruster usage, then, was during the reaction wheel biasing.

If the initial and desired final momentum states are known, the desired momentum change vector can be calculated in inertial space. Due to the configuration of the reaction wheels (see Fig. 2), a turn to align the spacecraft Z-axis with the desired momentum change vector results in an identical momentum change in each of the wheels. Because the wheels are equidistant from the spacecraft Z-axis, the bias will apply a torque purely about the spacecraft Z-axis - requiring only the Y-thrusters to fire.<sup>10</sup> For a more detailed description the y-bias design process, see Ref. 10.

Y-biasing on Cassini had the additional operational benefit of incurring very small  $\Delta V$  from thruster firing, due to the Y-thrusters being fired in equal and opposite pairs, unlike the Z-thrusters. These biases, therefore, caused smaller perturbations to the flight path and helped the spacecraft to fly closer to the reference trajectory.

After the swap to B-branch in early 2009, y-biases were used more often than other biasing methods. The Cassini mission ended with a more closely balanced total on-time between the B-branch Y- and Z-thrusters, as show in Table 2. Y-biases had an additional benefit in that they saved an average of 30-35% of bias hydrazine usage. They generally had a smaller duration than other biases, and therefore were more fuel-efficient.

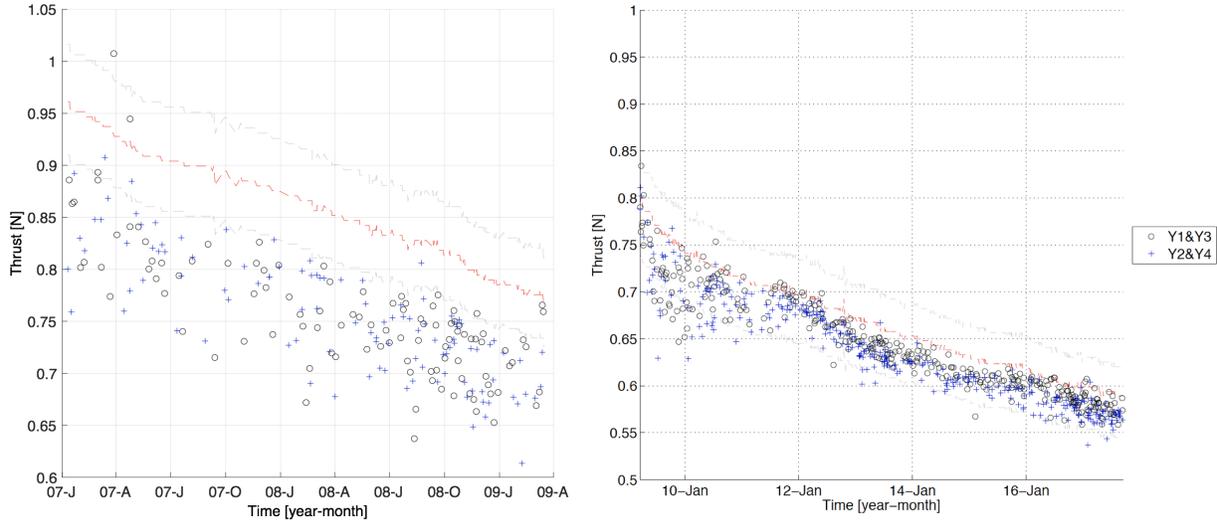


Figure 9. RCS Y-Thruster Branches Mission Performance.<sup>9</sup> A-Branch (left) and B-Branch (right) Y-Thrusters Pairs Thrust Magnitude Estimates. Minimum on-time of 5 seconds. Solid (red) line shows expected thrust. Dash-dot (black) lines show  $\pm 5\%$ , an envelope of expected deviation. Circles (black) indicate estimates for the Y1/Y3 pair while crosses (blue) represent the Y2/Y4 pair.

#### IV. Inertial Reference Units (Gyroscopes)

Cassini’s attitude estimation algorithm used data from both the IRU and the SRU. Cassini had two fully redundant IRUs – IRU-A and IRU-B. Each Cassini IRU contained four Hemispheric Resonator Gyroscopes (HRGs). Three of the HRGs were oriented orthogonally and used as the primary angular rate sensors. The fourth HRG was skew-mounted compared with the other three and was used as a parity checker.<sup>11</sup> The IRU electronics maintained a resonant vibration in each of the quartz gyros, and returned the measured frequencies. As the IRU assembly rotated, the vibration frequency in the gyros changed. These variations were correlated with the IRU assembly (and spacecraft) rotation rates.<sup>12</sup> The IRUs required periodic calibrations, as scale factor errors and misalignments caused the rate information from the IRUs to drift over time. This drift, even with calibrations, limited the time that a spacecraft attitude could be accurately propagated without updates from the star tracker. An overview of Cassini’s calibration strategy, along with trending of the scale factor and bias of the IMUs, is provided in Section IV. A.

IRU-A was powered on continuously from launch to the end of mission, meaning that IRU-A operated for more than 175,000 hours in flight. The back-up IRU, IRU-B, was only powered on for periodic calibrations as well as a few mission-critical spacecraft events, such as Saturn Orbit Insertion (SOI). During these critical events, IRU-B was powered on as a “hot backup” in the event that a spacecraft anomaly necessitated an immediate swap to the back-up hardware.

The IRUs did experience a small number of Single Event Transients (SETs), at which time radiation and increased solar activity caused spikes in a fault protection monitor that came close to triggered a spacecraft fault response. The overall nominal performance of the IRUs was not affected, and the spacecraft team raised the limits on this monitor to allow margin for future occurrences. This is discussed further in Section IV. B.

##### A. Hardware Performance

IRU-A was used continuously for Cassini’s Cruise from Earth to Saturn (1997–2004). However, the first calibration of IRU-A was not performed until 2002, while the first IRU-B calibration was performed in 2003. The results of these first calibrations led to IRU parameter updates in the flight software before the Saturn Orbit Insertion (SOI) burn. After arrival at Saturn and the start of the tour, Cassini’s IRU-A was calibrated every year and IRU-B was calibrated approximately every other year, through the end of the mission in 2017.

Each IRU calibration involved a change in the standard telemetry schedule and attitude estimator mode, followed by the execution of approximately  $180^\circ$  slews about each of the spacecraft body frame  $\pm X$ -,  $\pm Y$ -,

and  $\pm Z$ -axes. The calibration mode of the attitude estimator allowed automated on-board estimation of the scale factor errors and misalignments, in addition to the standard in-flight estimation of the IRU bias. The calibration slews had to be carefully constructed by an AACS ground operator to avoid turns that would point the SRU at a bright body (which could corrupt the inertial attitude knowledge in the estimator). Each calibration activity took between four and five hours. They did not consume any propellant, however, as the slews were all executed on reaction wheel control. For more detailed discussions of the IRU calibration procedure, see Refs. 12 and 13.

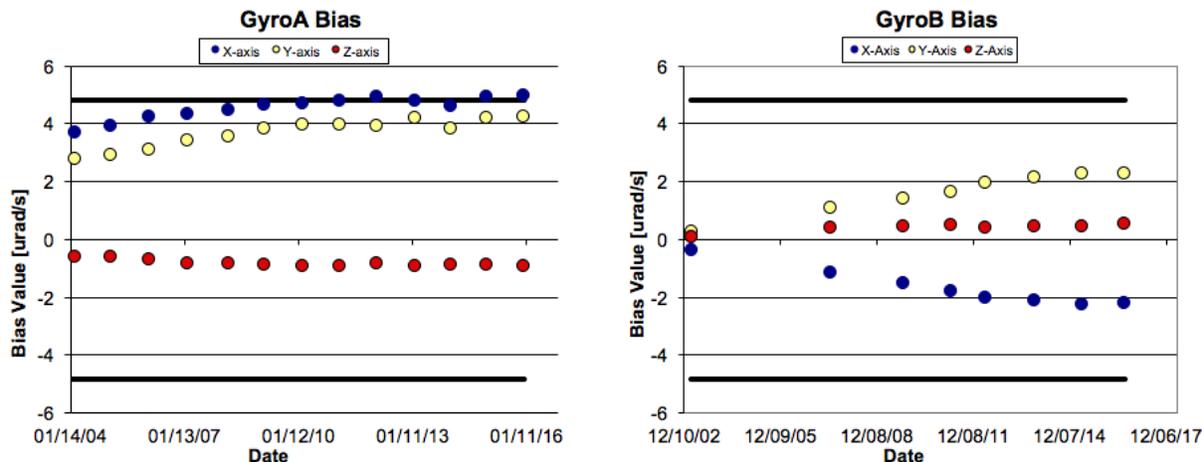


Figure 10. Cassini IRU Per-Axis Bias Trending

Fig. 10 shows the history of the Cassini IRU per-axis bias. The bold, black, horizontal lines represent the uncalibrated launch requirement of  $1^\circ/\text{hr}$  (or approximately  $4.85 \mu\text{rad/s}$ ). The Y- and Z-axes have remained below this requirement for the duration of the mission, while the X-axis bias drifted to slightly above the un-calibrated requirement at launch. This was not a concern for attitude estimation because the GNC software was continuously estimating the biases using a Kalman-Bucy filter.

The history of the IRU scale factor errors is shown in Fig. 11. The scale factor errors remained below 0.10% for the majority of the mission. The largest change in the scale factor over one year (not including the changes due to flight software updates) was approximately 0.03%, on IRU-A from 2008 to 2009. The requirement for scale factor stability, which required that the scale factor change no more than 0.05% over 30 days, was met. After the initial calibrations for both IRU-A and IRU-B showed large scale factor errors (0.2% on the IRU-A Y-axis), updates were made to scale factor errors for both sensors, designated in Fig. 11 by the vertical dotted lines. There was only one other update to scale factor errors, in 2006 on IRU-A, made for the rest of the mission. The scale factor errors of both for IRU-A and IRU-B remained relatively constant for the last 10 years of the mission.

The IRU misalignment history is given in Fig. 12. Like for the scale factor errors, updates were made to the misalignment parameters in the flight software for both IRU-A and IRU-B after the first calibrations in 2002 and 2003 (designated again by the vertical dotted lines). These were the only updates made to the misalignments for the entire mission. The misalignments varied very little through the years at Saturn. Overall, the Cassini IRUs performed nominally and required few changes to the on-board IRU parameters.

## B. Single Event Transients

Saturn’s radiation environment, while not strong enough to be cause for concern regarding hardware lifetime, did cause numerous Single Event Transients (SETs), also known as Single Event Upsets (SEUs), where a charged particle caused an erroneous, instantaneous spike in various measurements across many spacecraft subsystems. Depending on the area affected, the spacecraft response to the SET could range from mild to severe. These events had a loose correlation with ring-plane crossings close to one of Saturn’s rings, as the rings absorb many of the charged particles, but Cassini also saw SETs on the way to Saturn during Cruise.

Of the Cassini AACS hardware, the IRUs were the most affected by SETs. When a high energy proton

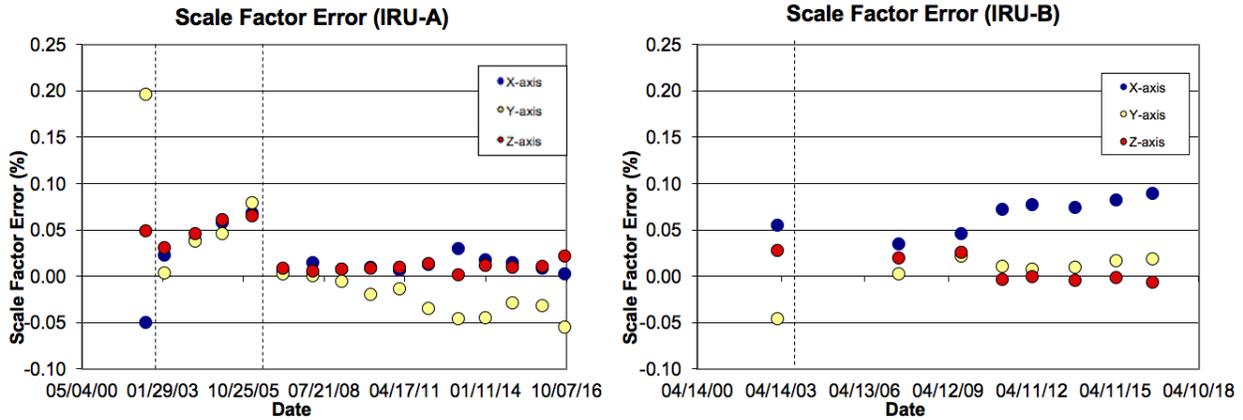


Figure 11. Cassini IRU Per-Axis Scale Factor Error

or heavy ion impacted the buffer circuitry of the gyroscopes, the IRUs would erroneously show a spike in the measured spacecraft rates. This would manifest as a small perturbation in the attitude control error (usually below  $50 \mu\text{rad}$ , which had a negligible impact on spacecraft pointing accuracy because the IRU data is filtered). However, an SET could also result in a large spike in one of the spacecraft fault monitors - the IRU parity check. This monitor combined the information from all four gyroscopes in one IRU into a sum parity that would nominally be very close to zero. A parity sum that was larger than a specified threshold for a duration longer than a specified persistence would trigger a fault response that assumed that there was an error in one of the four gyroscope axes.<sup>2</sup>

In the summer of 2000, it was observed that there had been several SET events when the IRU parity monitor exceeded the threshold and closely approached the persistence limit. To provide margin against an SET being misdiagnosed as a real problem with the prime IRU, the persistence limit was increased in August of 2000. There are more details of the change in persistence limit given in Ref.<sup>2</sup>

## V. Stellar Reference Units (Star Trackers)

Cassini had two redundant Stellar Reference Units (SRUs), or star trackers, SRU-A and SRU-B. The SRU was used, along with the IMU, to estimate the spacecraft attitude in the celestial frame. Cassini's attitude estimator continuously propagated the attitude with the body rate information provided by IRUs at a high rate, and ingested the attitude estimate from the star trackers each time a star picture was taken. Cassini's SRUs were charge coupled device (CCD) star trackers that were able to provide autonomous star identification by taking images and matching stars to an on-board star catalog. The optical field of view (FOV) of each SRU was  $\pm 7.5^\circ$  by  $\pm 7.5^\circ$ , and was  $1024 \times 1024$  pixels with a resolution of  $255.6 \mu\text{rad}$  per pixel.<sup>2</sup> The star catalog contained information about position, color, magnitude, and usability of stars used for star identification (or SID). An inertial attitude estimate was obtained by matching two to five stars in an image from the SRU to the star catalog. The SRUs were each calibrated approximately once a year to track misalignments between the two sensors, as well as "hot pixels" (which could interfere with the star identification) on the SRU CCD.

Overall, Cassini's star trackers performed well for the duration of the mission. There was one spacecraft anomaly in 1998, soon after launch, that resulted in a spacecraft safing response (see Section V. B) due to an unrealistically tight threshold on a fault protection monitor relating to attitude estimation (triggered by the small misalignment between the prime and back-up SRUs). The mission began with SRU-A as the prime star tracker, but during the course of this safing event, SRU-B became prime. As SRU-B continued to perform nominally, it remained the prime star tracker for the remainder of the mission. SRU-A was powered on for yearly calibrations.

The SRUs did require careful operations work to anticipate and manage bright bodies in the FOV of the star tracker. A sufficiently large bright body in the images taken by an SRU would prevent the sensor from producing an accurate attitude estimate.

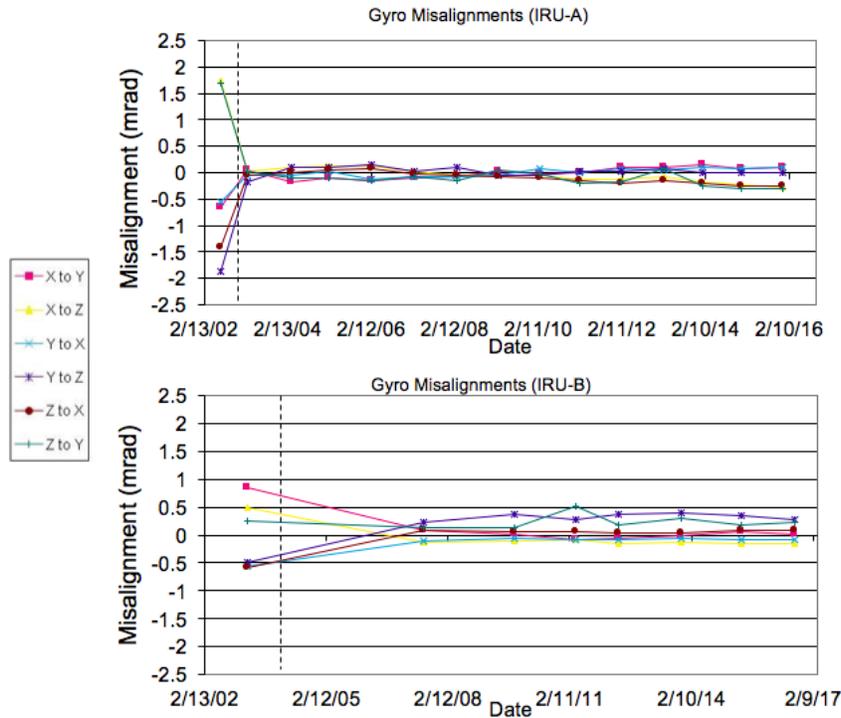


Figure 12. Cassini IRU Misalignments

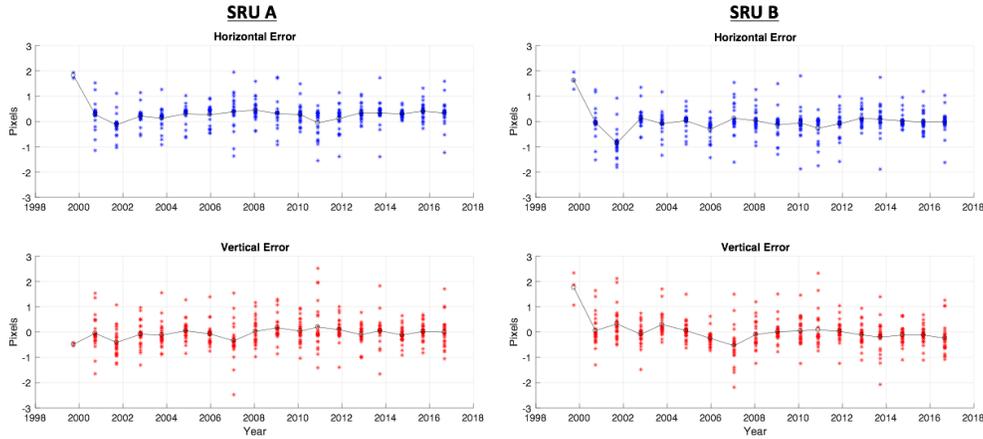
### A. Hardware Performance

Cassini launched with SRU-A as the prime star tracker. In March of 1998, the first calibration of the SRUs was performed to determine the misalignment between SRU-A and SRU-B. During this time, SRU-B was made the prime sensor for the first time. The swap resulted in a spacecraft anomaly, and SRU-B remained the prime sensor for the rest of the mission - but both star trackers continued to demonstrate nominal performance. The performance of the SRUs was tracked through calibrations that occurred one per year. The calibrations suspended the star tracking process, one at a time, on each star tracker to collect information about the location and intensity in SRU's field of view of the brightest 25 identified spots. On the ground, the star catalog was simulated using knowledge of the spacecraft attitude from telemetry. The downlinked SRU spot locations in the FOV were compared with the expected positions of cataloged stars.

The processing of such data on the ground could reveal misalignments, focal length errors, and "hot pixels". A misalignment would appear as a persistent offset in either or both the vertical and horizontal positions of the measured spot as compared to the matched star from the catalog. A nominal focal length of the star tracker was a nonlinear function of radial distance from the center of the SRU FOV. An error in focal length would appear as errors between measured spots and matched stars radiating from the center of the SRU FOV. A hot pixel was a small region of the star tracker's optical lens that was influenced by charged particles and would persistently and falsely identify as a star. In the calibrations, a hot pixel would have a centroid and intensity that was unable to be matched to any star in the star catalog.

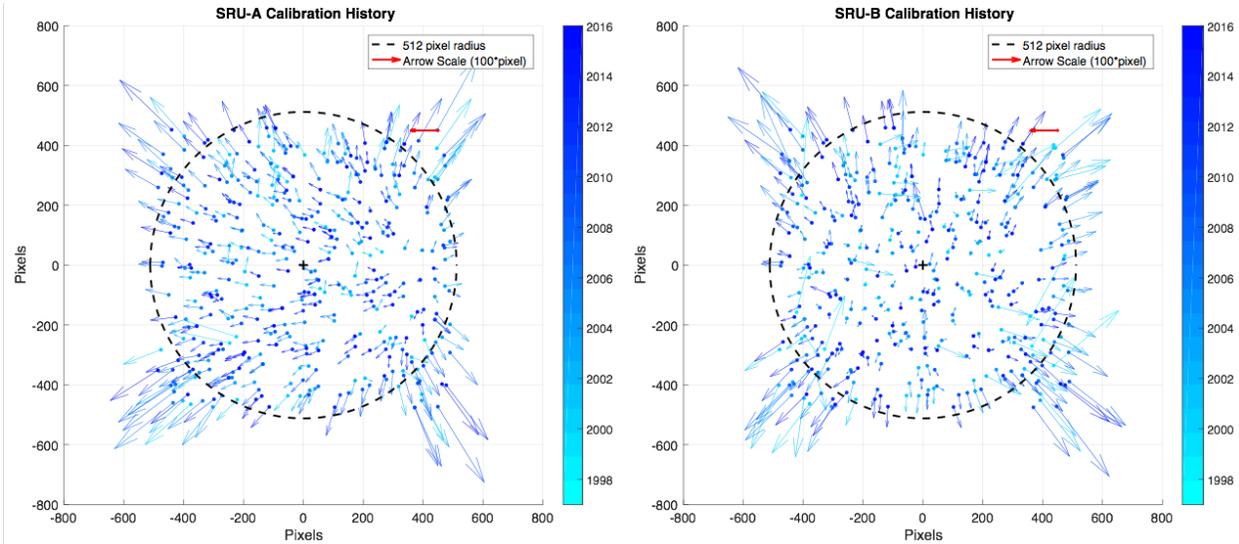
Fig. 13 shows the history of the comparison between the "stars" identified in the calibration and the expected star positions based on ground calculations using the downlinked spacecraft attitude and the star catalog. The very first calibration in March of 1998 resulted in the first spacecraft safe mode, discussed in Section V. B.

Due to the safing, there was no calibration data available for the calibration in March of 1998. The calibration in December of 1998 was unsuccessful due to a radiation-induced error with the Solid State Power Switches (SSPS) for the SRUs. There are therefore no stars shown for the calibrations in 1998. The calibration in 1999 resulted in very few stars, as seen in Fig. 13. The attitude of the spacecraft for this calibration pointed the SRU boresights to a region of the sky where there happened to be few stars. All



**Figure 13. Comparison Between Cassini SRU Calibration Telemetry Spots and StarID Catalog. Errors are given both in the horizontal SRU FOV axis and in the vertical SRU FOV axis.**

calibrations before 2002 were executed on RCS control - spacecraft motion from thruster deadbanding could slightly increase the error between the measured spots and the ground-matched stars during those years. The mean spot horizontal and vertical errors (represented by the black line in Fig. 13) for both trackers showed a very consistent trend for the calibrations in which many stars were available for comparison. Both mean SRU-B errors, horizontal and vertical, remained negligibly small for the duration of the mission. The mean SRU-A vertical error also was negligible. The mean SRU-A horizontal error was consistently just under 0.5 pixels, which translates to a slight misalignment of approximately  $100 \mu\text{rad}$ . As the Cassini inertial attitude knowledge requirement was 1 mrad (radial 99%),<sup>2</sup> the small misalignment in SRU-A was not of concern and therefore not corrected in a new version of flight software.



**Figure 14. History of Star Position Error in the SRU FOV**

In Fig. 14, the history of the SRU calibration results are shown in the frame of the SRU boresight. Each dot represents the expected position of the star based on the downlinked knowledge of the spacecraft attitude. The arrows represent the direction and the magnitude of the error. The magnitude of the arrows are multiplied by 100 to make them visible, so the true magnitude errors are two orders of magnitude smaller than shown in the figure. The red arrow shows the exaggerated arrow length that is equivalent to 1 pixel. In these figures, a focal length error would appear as all of the arrows pointing away from the center of the FOV. Overall, there does not appear to be a trend over time in either SRU. In the SRU-A plot of Fig. 14, the

large majority of arrows are pointing in the same direction - mostly horizontal. This is the same evidence of the slight horizontal misalignment in SRU-A that is visible in Fig. 13, but shown in a different way. There is, however, no indication of a substantial focal length error. The errors at the edge of the FOV are larger and pointed away from the center of the FOV, although the magnitudes have never been much more than 1 pixel. This was an indication of a small focal length error that did not require correcting in the flight software. The same small focal length error is visible at the corners of SRU-B. There are no consistent misalignment errors visible in SRU-B, as there was in SRU-A. There were no updates made to the focal length in the flight software for either tracker over the course of the mission.

The history of the SRU's unmatched spots is shown in Fig. 15. The spots represent the pixels that showed a measured star during the calibration that was not matched to a catalog star. In a single calibration, unmatched stars could be (1.) true stars that were not stored in the star catalog, (2.) charged particle hits at the time of the calibration, or (3.) pixels in the SRU FOV that persistently and falsely identified as a star (or hot pixels). Random charged particle hits would not reoccur at the same position in subsequent calibrations. Additionally, as the calibrations were done at different inertial attitudes, spots that are true (but not cataloged) stars would not reoccur in future calibrations. Hot pixels would, however, reoccur at the same location in the FOV of the star tracker, regardless of attitude.

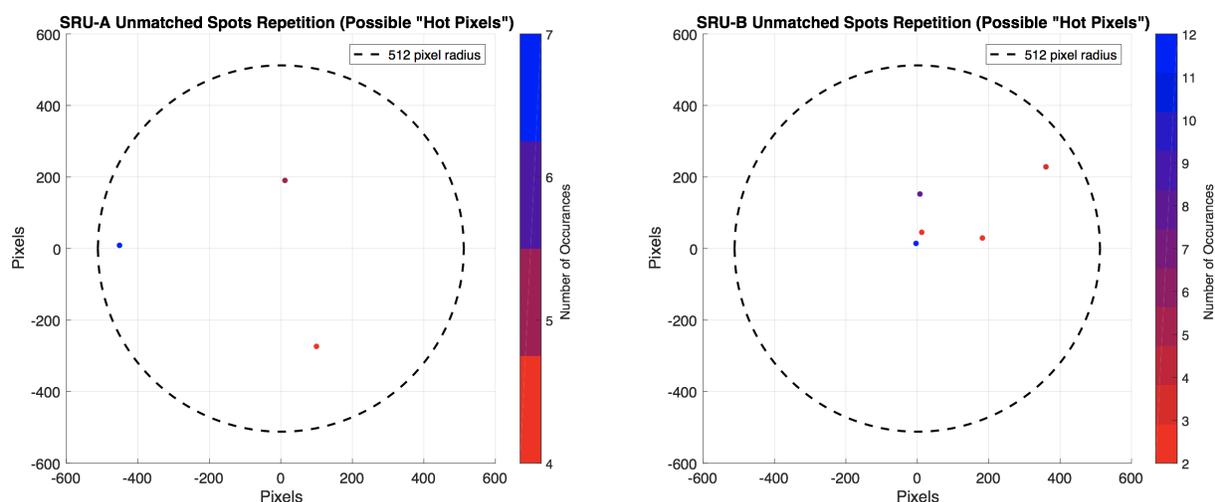


Figure 15. History of Repeated SRU Hot Pixels

Fig. 15 shows the history of persistent unmatched spots (the unmatched spots occurring only in one calibration were not included). The color represents the number of calibrations in which that unmatched spot appeared. As stated above, a hot pixel would appear at the same location in the star tracker FOV, regardless of spacecraft attitude. However, a calibration only stores the information for the brightest 25 stars. A calibration at an attitude with many bright stars may not capture a dimmer hot pixel, but that same hot pixel may appear in a subsequent calibration with fewer true, bright stars. Fortunately, these dim hot pixels were not a threat to the attitude estimator - it used a maximum of the five brightest stars in the SRU FOV. Only if a hot pixel was consistently bright enough to be used in the attitude estimator would there be potential for error in the spacecraft attitude. All of the observed hot pixels were less than half as bright as the brightest stars, and therefore were never considered a threat to spacecraft performance.

SRU calibrations also tracked the magnitude of intensity of each star. Early in the mission, it was observed that the intensity of the measured spots averaged approximately 1.5 times the star catalog values. In March of 2000, a parameter in the flight software was changed to bring the measured values closer to the catalog numbers.<sup>2</sup>

Both star trackers performed nominally for the duration of the Cassini mission. All hardware-related errors remained small and stable and did not require changes to the attitude control flight software to correct.

## B. Safe Mode March 24, 1998

Most spacecraft have software on-board called “fault protection” that monitors the spacecraft behavior and watches for anomalies. Fault protection software also contains instructions for a spacecraft to attempt to autonomously correct certain problems by, for example, switching to back-up hardware. There are some situations in which the best option is for a spacecraft to turn off all non-essential functions, turn towards Earth, and broadcast its status. This spacecraft mode is called “safe mode”. Faults can then be identified on the ground, and resolutions can be tested and uplinked to the waiting spacecraft.

Cassini entered into “safe mode” very few times compared with other missions. In the near 20 years of flight, the spacecraft entered into safe mode six times. Only one out of the six initiated safe modes were caused by faults related to the Cassini AACS (or GNC) hardware. The other initiated safe modes were caused by operator error, unexpected behavior of the flight system, and cosmic ray hits in critical pieces of non-AACS hardware.

Approximately five months after launch, on March 24, 1998, Cassini entered into safe mode for the first time. On-board, the spacecraft was undergoing the first SRU calibration. For the activity, SRU-B was powered on and made prime in the control loop for the first time. Less than two minutes after SRU-B became prime, fault protection called spacecraft safing in response to a large spike in the “Z-sigma ratio” fault monitor. Periodically, the propagated attitude estimate (based on the IRU and previous SRU attitude measurement) was compared with an SRU measurement update by computing a residual error vector. The “Z-sigma ratio” was the magnitude of this error residual vector, normalized by the error covariance matrix. The fault monitor triggered if the Z-sigma ratio ever exceeded a threshold stored in the flight software. Investigation revealed that, on March 24, 1998, there was a spike that exceeded the fault protection threshold in the Z-sigma ratio immediately after SRU-B was made prime. It was determined that the spike was caused by a misalignment between SRU-A and SRU-B. In response, Cassini operators increased the threshold of the Z-sigma ratio monitor in the flight software to be more robust to spikes in the future.

## C. Extended Bright Bodies

Before launch, it was known that there would be a need to suspend the star identification when there were sufficiently large bright bodies in the field of view of the star tracker. The supporting software development both on the ground and onboard the spacecraft was deferred until after the Jupiter flyby during Cruise, in December of 2000, to take advantage of flight data. During the Jupiter flyby, Jupiter always remained smaller than  $1^\circ$  in angular diameter relative to the spacecraft. Despite the small size, all star tracking was lost when Jupiter was in the FOV of the star tracker. The resulting perturbations in the attitude estimate were as large as 15 mrad.<sup>11</sup> Using the flight data from the Jupiter flyby, the star identification (SID) suspend protocol was developed. During an SID suspend, the attitude estimator would propagate the spacecraft attitude using only the information available from the IRUs. At the end of the suspend, SID was reenabled and attitude error introduced from the drift of the gyros was corrected. It was necessary to limit the duration of the SID suspends to avoid the need for excessive attitude error corrections, which would trigger fault protection monitors. The maximum duration of an SID suspend allowed in flight was five hours.

While the need for SID suspends was anticipated before launch, the need for careful management of Z-sigma monitor (discussed in Section V. B) was not expected. It was observed in flight data that there were some cases where a bright body in the FOV of the SRU did not require an SID suspend to avoid large attitude errors, but could trigger the Z-sigma monitor - the same monitor that caused the first spacecraft safe mode in 1998. While the threshold was increased after the anomaly in 1998, a small Saturnian moon in the FOV of the star tracker could still cause this monitor to be above the allowable flight threshold for long enough to cause another spacecraft safing. To avoid this occurrence, “Z-sigma masking” activities (where the fault protection monitor was masked, preventing any fault protection responses associated with that monitor) were developed and treated very similarly to the SID suspend activities.

For both SID suspends and Z-sigma maskings, the commands were developed on the ground during the development of each science sequence. A ground tool kinematically simulated the spacecraft attitude and identified the times when the geometry of bright bodies required an SID suspend or a Z-sigma mask. For a thorough discussion of the development of both SID suspends and Z-sigma masks, see Ref. 11.

## VI. Sun Sensors

The Cassini spacecraft had two redundant Sun sensor assemblies (SSA). The Sun sensor information was nominally not used in the control algorithms - the sensors were only nominally used in a fault monitor to compare the provided SSA Sun direction (when available) to the spacecraft's expected Sun position. The Sun sensors were most crucial in off-nominal scenarios, when a loss of inertial attitude knowledge was caused by a spacecraft fault and needed to be reacquired. The first step was to autonomously find the Sun by using the SSA and to achieve Sun-point. After that, 3-axis celestial attitude knowledge is acquired by using SRU star field images in star identification flight software. Both Sun sensors were mounted on the high-gain antenna (HGA) of the spacecraft (see Fig. 1), with the center of the FOVs aligned with the spacecraft -Z body axis.

The Sun sensors were digital Sun sensors, each with a square field of view of  $\pm 32^\circ$  by  $\pm 32^\circ$ . Each sensing axis consisted of a grid of small solar cell, housing with a narrow slit, and an associated electronics box. The light passing through the slit of the sensor was cast onto the solar cells, and produced a Gray code that flight software converted to an angle.

The prime Sun sensor assembly, SSA-A, was prime for the duration of the mission. SSA-B was turned on for Sun sensor checkouts (discussed in Section VI.A), and was powered on for the Grand Final (i.e. the last several months of the mission). There was no designated calibration cadence for the Sun sensors as there was with the other AACS sensors. The "unexpected Sun position" monitor, which triggered if the angle of the Sun from the Sun sensors deviated from the attitude controller estimate by more than a defined threshold, was regularly monitored. There was also a monitor that would trigger if the Sun was expected to be in the FOV of the Sun sensor, but the Sun sensors were not returning a reading. This monitor required careful handling of eclipses and was masked during any time when it was expected that the Sun sensor could not see the Sun to avoid a spacecraft fault response. The spacecraft did execute Sun sensor checkouts after spacecraft activities that could have endangered the Sun sensors - specifically, crossings of Saturn's ring plane near one of the rings.

### A. Hardware Performance

When the Sun was in the FOV of the Sun sensor and the SSA indicated that the Sun position estimate was valid, the "unexpected Sun position" monitor converted the attitude estimator's estimate of the Sun position into the coordinate frame of the Sun sensor, and compared the resulting angles to the Sun angles being measured by the Sun sensor. The threshold for this monitor was  $5^\circ$ . The value of the unexpected Sun position monitor was not stored and downlinked in telemetry, but the high water mark channel (HWM) was. A high water mark channel stored the maximum value that a corresponding fault protection monitor had seen. The HWMs were cleared every downlink pass to provide visibility into the behavior of the associated fault monitors. For the unexpected Sun position monitor, when either the hardware was powered off or the Sun was not in the FOV, the HWM was reported as the last valid value.

The histories of the unexpected Sun position HWM for both SSA-A and SSA-B are shown in Fig. 16. The resolution of the Sun sensor angles from the Sun sensors was coarse, only  $1^\circ$ . This is the reason for what appears to be instantaneous error jumps. SSA-B remained powered off for most of the mission - therefore, the HWM remained mostly at 0%. The occasional spikes shown for SSA-B are the times when it was powered on during a Sun sensor checkout for long enough for the unexpected Sun position to increase to  $1^\circ$  or  $2^\circ$ . As the HWM was cleared approximately every two days, the HWM value returned to 0% quickly each time there was a spike. The HWM was still reported even when SSA-B was powered off. SSA-A remained powered on for most of the mission. As seen in Fig. 16, the difference between the SSA Sun angles and the attitude estimator's Sun angles was consistently around  $2^\circ$  (40% of  $5^\circ$ ). This error includes mechanical misalignment, which was approximately  $1^\circ$ . Therefore, the error of the Sun position from the Sun sensors was  $1^\circ$  or less for most of Cassini's mission, which was less than the  $1.5^\circ$  requirement. There are five instances across all SSAs where the high water mark increased close to, or past, 100%. These instances were erroneous spikes due to data corruption. SSA-B was powered off on the days that the spikes in the SSA-B telemetry occurred and therefore could not have produced a Sun position to trigger the monitor. The sole spike on SSA-A (axis-2) was also erroneous - when the direct SSA-A telemetry was compared with the predicted Sun position during this time, the error angle remained well below  $3^\circ$ .

Occasionally during the mission, the spacecraft crossed Saturn's ring plane near one of the dusty rings. At these times, the spacecraft pointed the HGA in the direction of the on-coming dust to protect the rest of

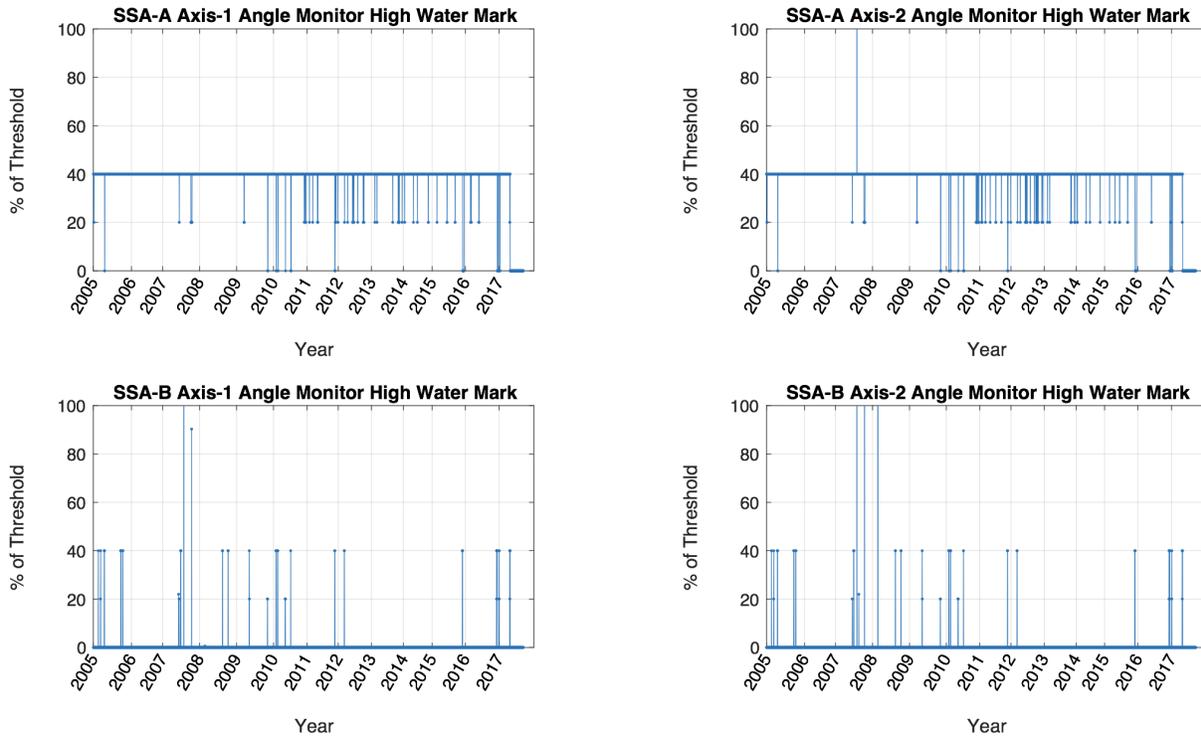


Figure 16. History of the Unexpected Sun Position Monitor High Water Mark

the spacecraft from micro-meteoroid impacts. The SSAs, however, were mounted on the HGA (see Fig. 1) and therefore were exposed to the on-coming particles. After each of these “hazardous ring plane crossings”, the spacecraft would power on the back-up SSA and execute a roll slightly off of the Sun-line. The telemetry was then compared on the ground with the predicted Sun position to verify that both SSAs were healthy. Fig. 17 shows an example of a Sun sensor checkout done on May 30, 2017. The blue line represents the ground-predicted position of the Sun, while the red and green lines represent the telemetry from both SSA-A and SSA-B.

The Sun sensors performed well for Cassini’s mission and did not require periodic calibrations to ensure the continued nominal performance of the sensors.

## VII. Accelerometer

Cassini had one single-axis accelerometer (ACC), aligned with the spacecraft +Z-axis. There was no back-up accelerometer - in the event of a failure, a “pseudo-accelerometer” (i.e. an equation in the flight software) would be used. The accelerometer was only used for  $\Delta V$  maneuvers that used the main engine, and spent the rest of the mission powered off. The resolution of the accelerometer was approximately 2.02 mm/s per DN ( $\pm 3\%$ ), which was too coarse to sufficiently measure small burns using the RCS thrusters.

The ACC bias was calibrated before every main engine maneuver. An hour before each burn, the accelerometer was powered on and collected data over a one minute span (that is, the calibration error is about  $2.02/60 = 0.034$  mm/s<sup>2</sup>, which is better than the 4- $\mu$ g requirement). The data was sent to the flight software and the bias calculated. If the calculated accelerometer bias falls within a pre-selected range, it is considered valid and will be used in the burn termination logic to cutoff the burn.

### A. Hardware Performance

Fig. 18 shows the history of the accelerometer bias, measured at each main engine  $\Delta V$  maneuver. The very first maneuver, TCM-1 (Trajectory Correction Maneuver #1), resulted in a burn 1.6% larger in magnitude than predicted. Much of this “overburn” was expected, due to a known error in the on-board accelerometer

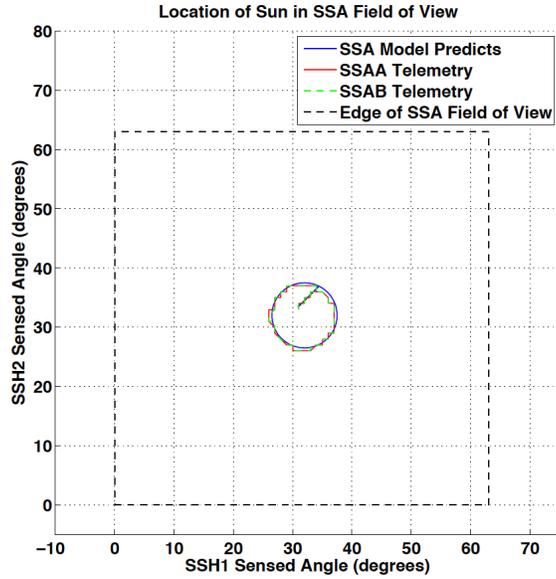


Figure 17. Sun Sensor Checkout after a Hazardous Ring Crossing on May 30, 2017

scale factor parameter. Before the next burn, the accelerometer scale factor parameter was corrected. This can be observed in the very first entry in the accelerometer bias plot in Fig. 18, which is an outlier compared with the rest of the accelerometer bias measurements. The bias change during the Cruise phase was relatively significant because of the large dynamic range (i.e. the change in distance from the Sun). During the prime mission, the bias change step is on the order of  $0.04 \text{ mm/s}^2$  consistent with the bias calibration error of  $0.034 \text{ mm/s}^2$  computed above. Over the mission, the accelerometer bias drifted from approximately  $3.07 \text{ mm/s}^2$  to as low as  $2.66 \text{ mm/s}^2$ . This is a total drift over the course of the mission of  $0.41 \text{ mm/s}^2$ . The requirement on the maximum drift over the whole mission was  $150 \mu\text{g}$ , or  $1.47 \text{ mm/s}^2$ . Therefore, the requirement on bias drift was met.

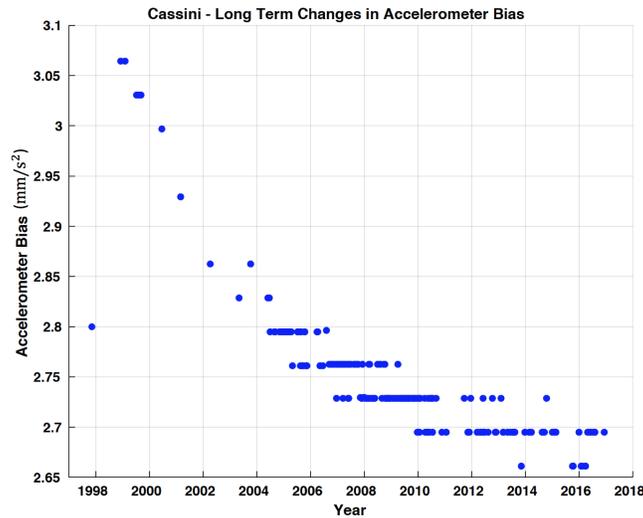


Figure 18. History of Accelerometer Bias

## VIII. Conclusion

The Cassini-Huygens mission was one of NASA's most successful missions. A well-designed spacecraft with robust hardware permitted two extended missions, lasting nine years past the end of the prime mission in 2008. The Attitude Control and Articulation Subsystem hardware performed nominally overall.

The AACS mission operations team placed a high priority on the ability to trend the performance of all GNC sensors and actuators throughout the mission. To allow for more efficient trending, several new ground software tools were developed in flight, such as the Thrust Estimation Tool. There was also new ground software developed in response to observed trends. For example, RBOT was developed to protect the health of the reaction wheels after RWA3 began to exhibit signs of cage instability. The trends were periodically reported to the AACS team and to project management. If anomalous trends were observed, the AACS team would also provide recommendations on resolutions. The resolutions ranged from being as simple as a change in the value of a fault monitor threshold to being as extreme as a thruster branch swap. The results of the trending tools were also frequently reported to members of the team who were involved in the pre-launch design of Cassini. These experts would provide an additional assessment of the trend information in an effort to identify anomalous trends early.

Ultimately, the lifetime of the Cassini spacecraft was not limited by hardware performance, but by the main engine bi-propellant. Planetary protection requirements necessitated Cassini's mission end in 2017 as the spacecraft depleted its reserves of propellant, and Cassini plunged into Saturn's atmosphere with a healthy Attitude and Articulation Control Subsystem.

## Acknowledgments

The work described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Thank you to Thomas A. Burk, Allan Y. Lee, Laura Burke, and Julie Webster, senior members of the Cassini operations team who reviewed the paper and supported the collection of information over the many years of the mission. Thank you also to Todd Brown for his work in providing the reaction wheel, IRU, and accelerometer trending.

## References

- <sup>1</sup>Gibbs, R., “Cassini Spacecraft Design,” Proc. SPIE Vol. 2803, p. 246-258, Cassini/Huygens: A Mission to the Saturnian Systems, Linda Horn; Editor.
- <sup>2</sup>Lee, A.Y. and Hanover, Gene, “Cassini Spacecraft Attitude Control System Flight Performance,” AIAA-2005-6269, *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, August 15–18, 2005, San Francisco, California.
- <sup>3</sup>Brown, T. S., “A Full Mission (1997-2017) of Attitude Control Performance during Orbit Trim Maneuvers Performed by the Cassini-Huygens Spacecraft,” *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Kissimmee, Florida, January 8–12, 2018.
- <sup>4</sup>Brown, T.S., “In-Flight Position Calibration of the Cassini Articulated Reaction Wheel Assembly,” AIAA 2012-4539, *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Minneapolis, Minnesota USA, August 13–16, 2012.
- <sup>5</sup>Lee, C. C. and Lee, A. Y., “Cassini Reaction Wheel Momentum Bias Optimization Tool,” AIAA 2005-6271, *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, San Francisco, California USA, August 15–18, 2005.
- <sup>6</sup>Lee, A. Y. and Wang, E. K., “In-Flight Performance of Cassini Reaction Wheel Bearing Drag in 1997-2013,” *Journal of Spacecraft and Rockets*, Vol. 52, No. 2, March-April 2015.
- <sup>7</sup>Luna, Michael E. and Bates, David M., “The Challenge of Implementing RWA Biases on the Cassini Spacecraft during Extended Mission Sequence Development,” AIAA 2009-5764, *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Chicago, IL, August 10–13, 2009.
- <sup>8</sup>Mizukami, M., Barber, T. J., Christodoulou, L. N., Guernsey, C. S., and Haney, W. A., “Cassini Spacecraft Reaction Control System Thrusters Flight Experience,” *JANNAF Propulsion Meeting*, Colorado Springs, Colorado, May 2010.
- <sup>9</sup>Stupik, J., “Thruster-Specific Force Estimation and Trending of Cassini Hydrazine Thrusters at Saturn,” AIAA 2016-2087, *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, San Diego, California USA, January 4–8, 2016.
- <sup>10</sup>Brown, T. S., “Y-Biasing: A New Operational Method for Cassini to Control Thruster Usage while Managing Reaction Wheel Momentum,” AIAA 2010-7560, *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Toronto, Ontario Canada, August 2–5, 2010.
- <sup>11</sup>Sung, T. S. and Burk, T. A. “Extended Bright Bodies – Flight and Ground Software Challenges on the Cassini Mission at Saturn,” AIAA 2016-2090, *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, San Diego, California USA, January 4–8, 2016.
- <sup>12</sup>Brown, T., “Inflight Performance of the Cassini Hemispherical Quartz Resonator Gyro Inertial Reference Units,” AIAA-2013-4630, *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Boston, Massachusetts, August 19–22, 2013.
- <sup>13</sup>Burrough, E.L. and Lee, A.Y., “In-flight Characterization of the Cassini Spacecraft’s Inertial Reference Units,” AIAA-2007-6340, *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, August 20–23, 2007, Hilton Head, South Carolina.