SOLAR DYNAMICS OBSERVATORY
LAUNCH AND COMMISSIONING

**James R. O'Donnell Jr., Ph.D., Kristin L. Bourkland, Oscar C. Hsu,
Kuo-Chia (Alice) Liu, Ph.D., Paul A. C. Mason, Ph.D.,
Wendy M. Morgenstern, Angela M. Russo, Scott R. Starin, Melissa F. Vess[*]**

The Solar Dynamics Observatory (SDO) was launched on February 11, 2010. Over the next three months, the spacecraft was raised from its launch orbit into its final geosynchronous orbit and its systems and instruments were tested and calibrated in preparation for its desired ten year science mission studying the Sun. A great deal of activity during this time involved the spacecraft attitude control system (ACS); testing control modes, calibrating sensors and actuators, and using the ACS to help commission the spacecraft instruments and to control the propulsion system as the spacecraft was maneuvered into its final orbit.

This paper will discuss the chronology of the SDO launch and commissioning, showing the ACS analysis work performed to diagnose propellant slosh transient and attitude oscillation anomalies that were seen during commissioning, and to determine how to overcome them. The simulations and tests devised to demonstrate correct operation of all onboard ACS modes and the activities in support of instrument calibration will be discussed and the final maneuver plan performed to bring SDO on station will be shown. In addition to detailing these commissioning and anomaly resolution activities, the unique set of tests performed to characterize SDO's on-orbit jitter performance will be discussed.

INTRODUCTION

The Solar Dynamics Observatory (SDO) was successfully launched and deployed from its Atlas V launch vehicle on February 11, 2010. Three months later, on May 16, 2010, the fully commissioned heliophysics laboratory was handed over to Space Systems Mission Operations to begin its science mission. SDO is an Explorer-class mission now operating in a geosynchronous orbit, sending data twenty four hours a day to a dedicated ground station in White Sands, New Mexico. It carries a suite of instruments designed to observe the Sun in multiple wavelengths at unprecedented resolution. The Atmospheric Imaging Assembly (AIA) includes four telescopes with 4096×4096 focal plane CCDs that can image the full solar disk in seven extreme ultraviolet and three ultraviolet-visible wavelengths. The Extreme Ultraviolet Variability Experiment (EVE) collects time-correlated data on the activity of the Sun's corona. The Helioseismic and Magnetic Imager (HMI) enables study of pressure waves moving through the body of the Sun.

---

[*] Aerospace Engineer, Attitude Control Systems Engineering Branch, NASA Goddard Space Flight Center, Mail Code 591, Greenbelt, Maryland 20771, USA.

OVERVIEW OF THE SOLAR DYNAMICS OBSERVATORY MISSION

Figure 1 shows an image the SDO spacecraft. The side of the spacecraft shown at the top of the figure is its Sun-pointing side. The overall length of SDO along the Sun-pointing axis is 4.5 m and each side of the spacecraft bus is 2.2 m. The span of the deployed solar arrays is 6.25 m. SDO launched with a total mass of 3000 kg, 1400 kg of which was propellant mass.

The science goal of SDO is to provide data to help understand the Sun's magnetic changes, to determine how its magnetic field is generated and structured and how that energy is released. SDO data and analysis will help develop the ability to predict the solar variations and their impact on the Earth. SDO's three instruments measure the properties of the Sun and solar activity. The large number of images and other measurements taken continuously over long periods of time will allow scientists to see the oscillations of the Sun. These patterns can be used to look into and through the Sun. More information about the SDO spacecraft and science mission can be found on the NASA SDO website.[*]


Figure 1: Solar Dynamics Observatory

---

[*] http://sdo.gsfc.nasa.gov

Attitude Control System Sensor and Actuator Suite

Figure 2 shows a mechanical drawing of the spacecraft giving the locations of the attitude control system (ACS) sensors and actuators and science instruments. The SDO ACS was designed to tolerate any single hardware fault and retain the capability to meet all requirements for science data quality. The suite of ACS sensors, actuators, and computational capabilities were selected and arranged for performance and maximal redundancy; SDO ACS failure detection and correction (FDC) depends to a large extent on hardware redundancy. More information on the SDO hardware placement can be found in Reference 1 and an extensive description of the SDO ACS failure detection and correction design can be found in Reference 2.


Figure 2: SDO Attitude Control System Hardware

ACS Sensors. The SDO sensor suite consists of sixteen Adcole coarse sun sensors (CSS), one Adcole digital sun sensor (DSS), two Galileo Avionica autonomous, quaternion-output star trackers (ST), and three Kearfott Two-Axis Rate Assemblies used as SDO's inertial reference units (IRUs). The CSSs are the only attitude sensors required in the most basic Sun-pointing mode. The sixteen CSSs are divided into two independent sets of eight sensors each (CSSA and CSSB), and each set of eight can provide an adequate Sun vector with any seven sensors being functional.

For fine attitude determination, an on-board Kalman filter provides adequate attitude knowledge with input from any two of the three fine-pointing units—DSS, ST1, and ST2. To avoid simultaneous blockages of both STs, they are mounted nearly perpendicular to the SDO Sun-

pointing axis (X axis), and far enough apart from each other that the Earth and Moon do not block both at the same time throughout the science collection phase of the mission. The IRUs are arranged so that the sensitive axes from two units are aligned with each of the three body axes of the Observatory. Thus, any two out of three IRUs will provide full three-axis rate information.

In addition to these sensors, the ACS also makes extensive use of the guide telescopes (GT) mounted as part of the AIA instrument. Because of the high accuracy of the SDO science instruments, the ACS uses the GT data as the best available knowledge of the Sun center. There are four GTs, with one mounted to each of the four science telescopes; the ACS only needs accurate information from one of the four GTs, selected by SDO scientists as the controlling guide telescope (CGT), to perform its science control duties. Each GT has a field of view (FOV) of 0.5 deg within which sunlight illuminates at least one photodiode, and polarity of the control signal is determined; this is called the acquisition range. When the Sun center is within approximately 90 arcseconds of the FOV center, the GT is capable of providing attitude information relative to the Sun vector accurate to about 2 arcseconds; this is referred to as the GT linear range and is required for accurate science data collection.

ACS Actuators. SDO guidance functions are actuated by four Goodrich 70-Nms reaction wheel assemblies (RWA) and eight Ampac 5-lbf attitude control thrusters. (While not used for attitude control, there is also one Aerojet R-4D model bi-propellant main engine producing 110 lbf of thrust used for orbit raising ΔV.) The RWAs are arranged in a pyramidal structure so that any set of three provides full three-axis control capability. The ACS thrusters are grouped into four pairs of thrusters, with one thruster of each pair linked to fuel and oxidizer (monomethyl hydrazine and nitrogen tetroxide) by independent manifolds. In this way, the catastrophic failure of any one thruster can only require the closing of one manifold, leaving the other set of four capable of performing all necessary ACS tasks.

In addition to attitude control activities, the ACS is also responsible for the pointing and control of the two high-gain antennas (HGA). Each antenna consists of a dish mounted on an elevation gimbal, with that mounted on an azimuth gimbal. If either antenna fails due to a gimbal or electrical failure, the mission can still be completed. The greatest danger posed by the HGAs is the irradiation of the spacecraft itself; HGA FDC provides protection from that event.

Electronics and Microprocessors. A full copy of the command and data handling system of the SDO spacecraft resides on each of two independent main processors (MP). Only one of these MPs is in control of the spacecraft at any time. Each of these MPs also operates an independent copy of the attitude control task (ACT) and the on-board ephemeris. Most flight software tasks, including the ACT, operate on a 200-ms cycle, except that the ephemeris operates on a one-second cycle. Nominally, it is the job of the ACT to sample sensor data and issue actuator commands to maintain attitude control of the spacecraft throughout the mission. There are also separate microprocessors that reside in each of two independent but cross-strapped attitude control electronics (ACE) boxes. These ACEs are always powered, but only one can be in control at any time. Both ACEs route data from the CSSs, RWAs, IRUs, and various pressure and temperature sensors to the MPs. Only the ACE in control accepts, validates, and passes through actuator commands to the RWAs and the propulsion system's valves and thrusters from the MP in control. If there is a disruption of ACT control over the ACE in control, that ACE will cease passing through ACT actuator commands and will instead begin issuing commands of its own to the reaction wheels from the Safehold mode running on it. Safehold is a simplified mode, being dependent only on the CSSs and the RWA tachometer readings (during eclipse, if IRU signals are available, they are used to null rates). If the primary ACE itself is disrupted, such as from a restart due to a single-event upset (SEU), the other ACE will detect this state and assume control.

Attitude Control System Mode Design

Figure 3 shows a diagram of the SDO ACS control modes and allowed transitions. The ACS has four RWA-actuated modes and two thruster-actuated modes. More details about the ACS in general and the control modes in particular can be found in References 1, 3, 4, and 5. As discussed in the previous section, one RWA-actuated mode resides on the ACE microprocessors; this mode is called Safehold. The other five modes reside in the ACT. Sun Acquisition (SunAcq) Mode performs an attitude function similar to Safehold, in that it simply maintains a power-positive, safe attitude with respect to the Sun using CSS signals. It differs from Safehold in that IRU signals are used for angular rate information at all times.

For all other modes, attitude determination (AD) is performed with some combination of the fine attitude sensors and propagation of IRU-derived rate information. An attitude solution may be initialized either by accepting a valid ST quaternion (nominal) or by uploading an estimate by ground command (available for testing and contingency). Once a solution is available, it may simply be propagated using rate sensors, as is always done in the thruster based modes, or it may be replaced either using one preferred ST or by ground override command. The most accurate solution is obtained by combining all available fine attitude data from the two STs, the DSS, and the IRUs using a Kalman filter.

Whatever AD method is selected in the software, Inertial Mode uses the solution for attitude error calculation against the target attitude in all three axes. Inertial has two sub-modes that differ only in the target calculation. One tracks a Sun-referenced target quaternion using the on-board ephemeris to predict the appropriate inertially referenced quaternion for the Sun-referenced state. The other maintains a commanded, absolute, inertially referenced quaternion. Science Mode, during which most science data are collected, uses one of the specialized GTs to point a commanded science reference boresight (SRB) accurately at the Sun. The roll error about that SRB is calculated using the same methods as Inertial, except that the target is always Sun-referenced.

The thruster modes are called DeltaH Mode and DeltaV Mode. DeltaH is used to manage system angular momentum. With no magnetic torquers to gradually dump momentum, the thrusters must be used occasionally to remove momentum. To maximize time between uses of DeltaH, the mode allows a non-zero angular momentum to be placed into the body, which can be set to the opposite of any predicted angular momentum change. The attitude target for DeltaH is simply the attitude estimate at mode entry. DeltaV is used for changing or maintaining orbit parameters. It uses an absolute, inertially referenced target similar to Inertial's absolute targeting, and that target may be updated by command during a DeltaV maneuver.

Some transitions between modes are not allowed. By placing the in-control ACE into Safehold mode, the ACS mode running on the MP is ignored, so Safehold may be reached from any MP mode. Any MP mode may transition to SunAcq or to Inertial, including self-transitions. Science mode is the only other mode that may self-transition, and it may also be entered autonomously from Inertial mode when the Sun is in the field-of-view of the controlling guide telescope. DeltaH may be entered from SunAcq or Inertial mode. However, Science and DeltaV may only be entered from Inertial mode, with Science accessible only when Sun-referenced targeting is active and DeltaV accessible only when absolute targeting is active. These restrictions avoid large attitude changes occurring due only to misunderstandings of the two targeting sub-modes in Inertial. Thrusters are always disabled upon exiting DeltaH or DeltaV modes.
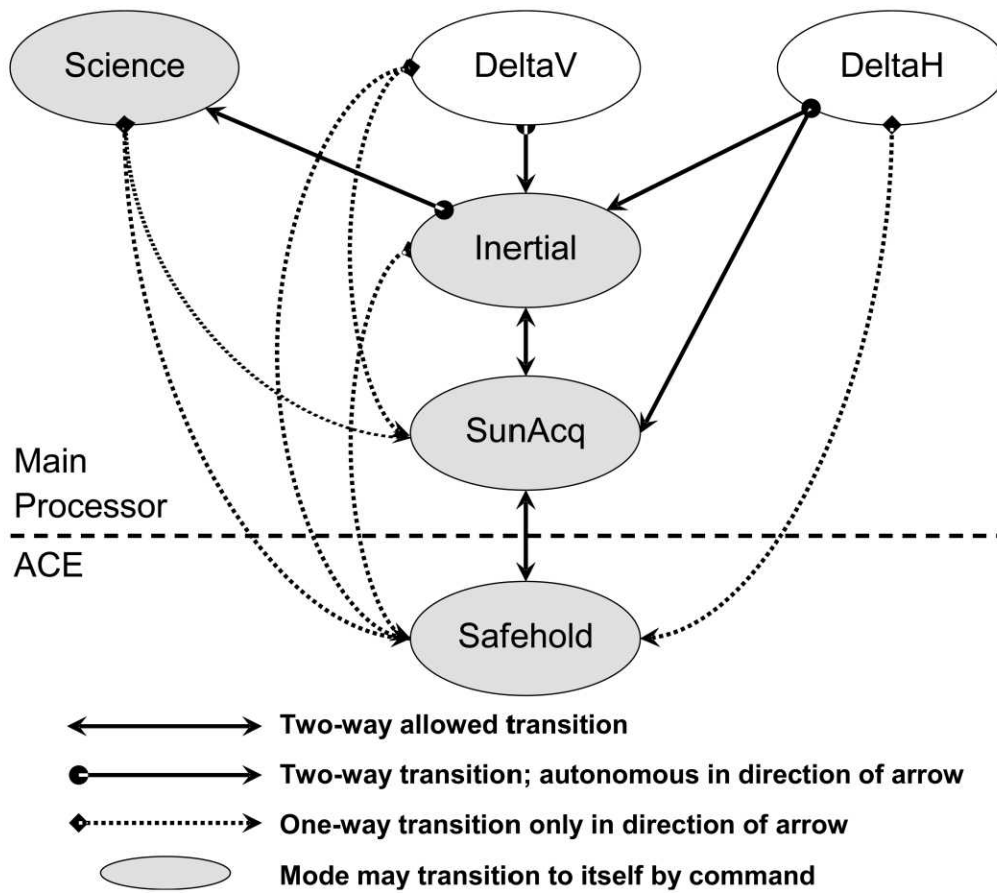
Figure 3: SDO Mode Transition Diagram

LAUNCH AND INITIAL ACQUISITION

After a one day launch slip caused by high winds at Cape Canaveral, SDO successfully launched at 15:23 GMT (10:23 EST) on February 11, 2010. On schedule an hour and thirty-six minutes later, the SDO Mission Operations Center established first contact with the spacecraft. At 17:07:39, SDO separated from the third stage of the Atlas V launch vehicle, having been safely delivered into its 8,800 km × 41,700 km, 28° geostationary transfer orbit. Figure 4 shows the SDO spacecraft rates leading up to and after spacecraft separation. (All times in plots are GMT.) The angular rates measured prior to separation are due to the third stage of the launch vehicle maneuvering SDO into its separation attitude. As shown, the separation event itself is relatively benign, imparting fairly low rates on the spacecraft, a total of 3 Nms of system momentum. SDO's solar arrays deployed automatically approximately five seconds after separation. Figure 5 shows the Sun angle, as measured by the CSSs, during initial acquisition. The requirement at initial acquisition for both SDO's SunAcq and Safehold Modes was to bring the spacecraft within 15° of the Sun within 30 minutes. In operation, SDO's SunAcq Mode was able to bring the spacecraft that close to the Sun within 15 seconds, with a total settling time of about three minutes. Note that, since the CSSs are located on the solar array panels, the Sun angles were not accurate until the solar arrays were deployed. Eighty-four minutes after separation, SDO's two HGAs were deployed. Figure 6 shows the spacecraft rate disturbance caused as each antenna was deployed; as expected, the primary disturbance is in the Y axis of the spacecraft.

Figure 4: Separation and Acquisition Spacecraft Rates

Figure 5: Post-Separation CSS Measured Sun Angle

Figure 6: Spacecraft Rates at High Gain Antenna Deploy

Safehold Checkout

One of the first post-separation and initial acquisition activities planned for SDO was a checkout of SDO's Safehold Mode. As the lowest level, minimum hardware component mode meant to ensure the safety of the spacecraft in the event of an anomaly, it was important to verify that the mode would work as designed. Additionally, because SDO had redundant ACE boxes in which Safehold was hosted, Safehold checkout was run on both ACEA and ACEB. The ACEA Safehold Mode test was run for ten minutes beginning at 21:50 GMT, followed immediately by a ten-minute test on ACEB. The plots shown in Figure 7 and Figure 8 are for the checkout of the ACEA Safehold Mode; ACEB yielded similar results.

SDO's Safehold Mode was designed to operate with a minimum sensor complement. This means that the only sensor inputs that it normally uses are the output from the CSSs and from the RWA tachometers. Using just the CSS inputs, Safehold Mode is able to determine its position with respect to the Sun and is able to derive its rate about the two directions transverse to the sunline. When the spacecraft is close to Sun-pointing (within 15 degrees), Safehold Mode determines the X axis rate—the rate about the sunline—from the changing momentum in the reaction wheels. If the RWA momentum is too low, the derived rate is deemed unreliable and the X axis rate error in the Safehold controller is zeroed. This limit is 2 Nms for the square root of the magnitude of the Y and Z components of the RWA momentum in the body frame.

During Safehold checkout, the spacecraft saw sunline rates which cycled between high and low values. The spacecraft remained Sun-pointing, but the roll around the sunline alternated be-

tween periods of low roll rates and higher roll rates. This response is evident in the RWA momentum in the body frame plot shown in the upper section of Figure 7.


Figure 7: Safehold Checkout RWA Momentum in the Body Frame

The lower plot of Figure 7 shows what is happening in this case. When the YZ component of the RWA momentum is below 2 Nms, the sunline error rate fed to the Safehold controller is zeroed. By zeroing out the small rate, the true error in RWA momentum is allowed to build up without being counteracted by the controller. At some point, though, the accumulated momentum in the YZ plane exceeds 2 Nms and the controller begins reacting to it. As the controller reacts to the error and begins to lower it, the total YZ momentum will again go under 2 Nms. This reaction is what causes the alternating periods of lower and higher sunline rates during the Safehold test. Figure 8 shows the corresponding sunline rate signal used by the controller. At a higher overall system momentum, the sunline rate estimation algorithm would behave more consistently and the overall control would be better. At the low system momentum that the checkout test was run, which is more typical of SDO's operations, the Safehold Mode rate estimation and control algorithms are sufficient to safely and stably keep the spacecraft pointed at the Sun, which is power and thermally safe, as designed.

After checkout of Safehold Mode, the final "initial checkout" activity performed was polarity checks of the two sets of ACS thruster command paths. These polarity checks were performed while the spacecraft was in SunAcq Mode via manual thruster pulse commands sent to the ACE, with the results of each thruster firing being evaluated by the induced spacecraft rate and change in system momentum. The checkout was performed through both the ACEA and ACEB boxes, with enough 200-ms pulses fired on each thruster to get consistent results from one firing to the next, in an effort to blow out any bubbles that might have gotten into the propellant lines. The results of these tests showed nominal operation of all ACS thrusters through both ACE boxes.

9

The checkout activities performed in the first day after SDO's launch and initial acquisition were done to verify that the spacecraft was in a good state of health and ready to begin its commissioning period, which would carry it through the next several months. The first major activity that was conducted during this period was the commissioning of SDO's hardware, software and science instruments. Additionally, a series of orbit-raising maneuvers were conducted in order to move SDO from its GTO orbit into its final mission orbit, a geosynchronous orbit at 102° west longitude, inclined at 28°. While these activities were conducted in parallel during the commissioning period, this paper will first discuss SDO's orbit raising, followed by the commissioning activities for SDO's ACS sensors, actuators, control modes, and software. Finally, a short description of the innovative tests performed to characterize SDO's on-orbit jitter performance will be shown.


Figure 8: Safehold Checkout Calculated Sunline Error Rate

ORBIT RAISING OF THE SOLAR DYNAMICS OBSERVATORY

The initial plan to carry SDO from its initial GTO of 8,800 km × 41,700 km to its final geosynchronous 42,164 km circular orbit consisted of ten maneuvers conducted over a period of three weeks. Table 1 shows the nominal orbit raising plan, consisting of one engineering burn, six apogee motor firings (AMFs), where the primary ΔV would come from SDO's 100 lbf main engine, followed by three trim burns (TMFs) performed using ACS thrusters only. The engineering burn was designed as a short "dress rehearsal" for a full apogee motor firing. In the engineering burn and each of the AMFs, the first 20 seconds of the burn would use ACS thrusters only, with thrusters off-pulsing for attitude control, and was meant as a settling burn to get the fuel settled in the propellant and oxidizer tanks before the main engine kicked in. After this settling burn, the rest of the burn would include the main engine, with the ACS thrusters in an on-pulsing mode for attitude control. See Reference 6 for a more complete description of the SDO orbit raising plan.

Table 1: SDO Nominal Orbit Raising Plan

Maneuver	Date and Time (GMT)	Burn Length (sec)	ΔV (m/s)
Engineering Burn	February 11, 19:19	80.0	10.67
Apogee Motor Firing 1	February 14, 02:55	1119.8	191.38
Apogee Motor Firing 2	February 16, 02:51	1040.0	188.74
Apogee Motor Firing 3	February 18, 06:48	1757.8	350.06
Apogee Motor Firing 4	February 20, 06:01	1126.8	245.77
Apogee Motor Firing 5	February 22, 13:47	607.2	139.18
Apogee Motor Firing 6	February 25, 04:01	533.4	127.12
Trim Motor Firing 1	February "28, 01:52	168.0	7.46
Trim Motor Firing 2	March 2, 12:21	193.6	8.62
Trim Motor Firing 3	March 4, 23:26	298.6	13.36

Propulsion System Commissioning and AMF-1

Due to heating concerns from one of the science instruments, the engineering burn was delayed from the nominal plan for four days, and was executed on February 15, at 02:42 GMT. Thruster polarity and rough thruster performance had been checked on launch day, so the only other propulsion system commissioning activity necessary before the engineering burn was to fire the pyrotechnic valves to pressurize the system. As shown in Figure 9, it was possible to see that the pyros had been fired based on disturbances to the spacecraft's attitude and rate.


Figure 9: Attitude Error Induced by Firing Propulsion System Pyros

The ACS performance during the engineering burn was as expected. There were some minor FDC limit threshold crossings that occurred during the burn, but nothing that raised any concerns. One such FDC indication was a derived DSS rate vs IRU rate comparison check that showed some failures that were eventually tracked down to error in the DSS rate derivation that occurred near the edges of its field of view. A second FDC indication seen, a failure of the comparison of the Sun angle as calculated by the A and B side CSSs and the attitude determination (AD) solution, will be discussed in the next section of this paper. As a result of the nominal performance of the engineering burn, the first apogee motor firing (AMF-1) was planned for February 17.

One day before AMF-1, the team elected to use manual thruster commands (firing thrusters in pairs) to reduce the system momentum in preparation for the upcoming burn. This was done manually rather than with SDO's DeltaH Mode for two reasons. First, the DeltaH Mode had not been tested yet and the team did not want any unexpected anomalies with that mode to delay the first full maneuver. Second, it was possible using these manual commands along with estimates of the daily system momentum buildup to bias the final momentum at the time of the unload such that by the time of AMF-1 the system momentum would be very low.

AMF-1 was executed at on February 17, with the burn starting at 22:15:28 GMT. The total length of the burn was 1133.4 seconds, slightly longer than the prelaunch plan, and completed successfully. The system momentum magnitude upon exit from DeltaV Mode was 3.07 Nms. Similar to the engineering burn, there were no failures or major problems during the maneuver, though some of the same FDC limit threshold crossings were seen (DSS vs IRU rate comparison, CSSA vs CSSB vs AD attitude comparison). The only thing that occurred during the maneuver of potential serious concern was the fact that the KA-band electronic box temperature hit its red limit just before the burn finished. As a result of this, the flight dynamics team was directed to replan the remaining AMFs to limit the burn time to 20 minutes or less, and new operational constraints and procedures were defined for future AMFs in case the thermal limits were reached again.

Figure 10 shows the attitude error during AMF-1, beginning one minute before the burn started. As expected, because the main engine is aligned along the X axis, the primary disturbances are in the Y and Z axes. The attitude errors during the maneuver are kept to within ±2°, for the most part, well within the ±5° requirement and FDC limit for DeltaV Mode. As also shown in Figure 10, the limit cycle during DeltaV Mode is one-sided for most of the burn. Towards the end of the burn, though, we see what looks like a two-sided limit cycle. Figure 11 shows the phase plane plot of the maneuver, as well as a plot of the system momentum magnitude. To give some sense of directionality on the phase plane plots, the data is plotted with the first third of the burn in red, the second third in green, and the last third in blue. As would be expected in the phase plane after seeing the attitude error, all of the limit cycles appear to be stable, with the Z axis being the largest. The system momentum magnitude plot is what would be hoped for in a well-behaved control mode. After an initial transient when the main engine kicks in after the settling burn (which is well within the FDC limit of 20 Nms), it varies for the rest of the maneuver at a fairly low level.

Based on the nominal real-time results of this maneuver, along with post-maneuver reconstruction from the flight dynamics team indicating expected ΔV performance from the thrusters, the planning process began for AMF-2, scheduled for two days later on February 19.

Propellant Slosh Anomaly, Investigation, and Recovery

AMF-2 was planned to be a 1080 second burn beginning on February 19 at 21:55:32 GMT. The burn began on schedule, but approximately 37 seconds into the burn, the high system momentum FDC limit was tripped, aborting the burn and putting the spacecraft into SunAcq Mode.
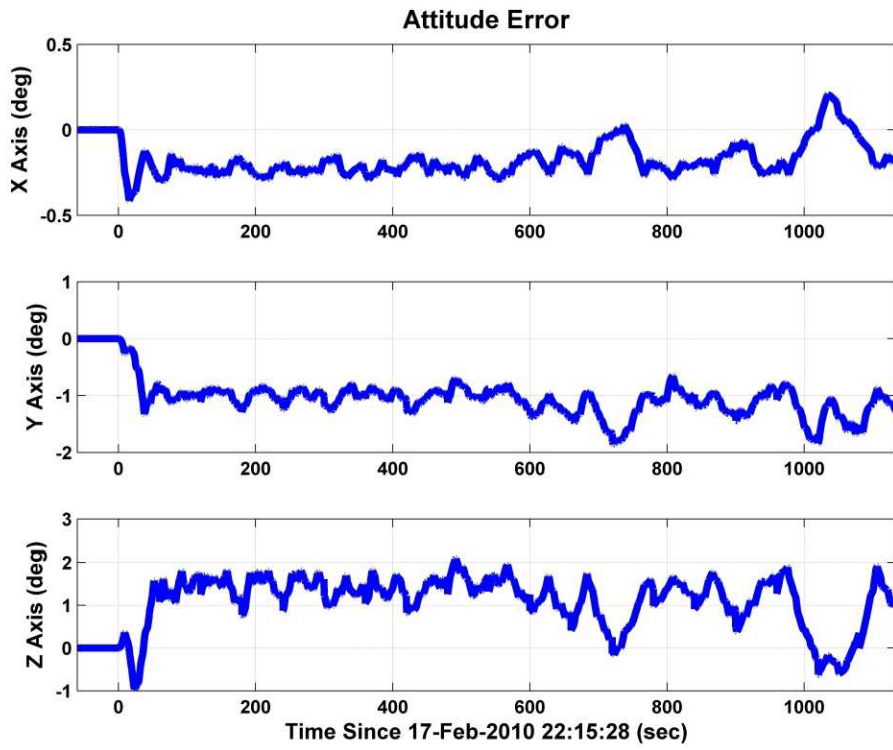
Attitude Error


Figure 10: AMF-1 Attitude Error

Figure 11: AMF-1 Phase Plane and System Momentum Magnitude

Once the FDC tripped and put the spacecraft into SunAcq Mode, that mode behaved as designed and quickly acquired the Sun. The calculated system momentum, while high enough at around 23 Nms to trip the FDC limit during DeltaV Mode, was well within the capability of SunAcq. The first action of the flight support and maneuver teams after the anomaly was to ensure the safety of the spacecraft and clean up the state of the spacecraft telemetry status monitors (TSMs) and other FDC components to protect the spacecraft while in SunAcq. Additionally, in order to reduce the relatively high momentum state of the spacecraft, manual thruster commands were performed to reduce the system momentum towards zero. Manual thruster firings were used to reduce this system momentum, with an as-expected momentum change occurring for each firing. As a first step towards debugging the anomaly, additional firings were performed using each thruster pair to verify correct operation of all of the ACS thrusters. After the spacecraft was safe, an anomaly team was formed to determine the cause of the DeltaV Mode anomaly.

Figure 12 shows the components and magnitude of the spacecraft system momentum during AMF-2. The momentum state upon beginning the maneuver (T = 0 on the plot) was relatively low at a magnitude of just under 5 Nms. When the main engine kicked in (T = 20), a large oscillatory change in the Z axis system momentum begins, and it is the momentum in this axis that drives the system momentum magnitude to exceed the FDC limit of 20 Nms and abort the burn after 37 seconds.


Figure 12: AMF-2 System Momentum and FDC Limit

The anomaly team identified a number of potential causes for the anomaly, including sensor failure, actuator (thruster or main engine) failure or previously unobserved misalignment, DeltaV Mode controller instability, other hardware failures, and unmodeled dynamics. Because of all the previous testing and examination of current data, there were not any sensor, actuator, or other hardware failures that seemed to be likely causes, and given its correct functioning during the

engineering burn and AMF-1, DeltaV controller instability did not seem likely. After examining the system momentum and rate data during the aborted maneuver, it quickly became clear that the most likely cause of the anomaly was an unmodeled disturbance caused by propellant slosh.

Once the anomaly team successfully identified propellant slosh as the reason the burn was aborted, they began an intensive analysis and design effort to come up with a way to mitigate the slosh dynamics to allow the use of the main engine. Meanwhile, because SDO was in an unfavorable radiation environment and because the desire was to get the spacecraft into its science orbit as quickly as possible, the team decided to continue with its orbit raising schedule, planning maneuvers using only the ACS thrusters. Because the ACS thrusters are much smaller and therefore less efficient, achieving the final science orbit without going back to the main engine would require an additional five or six burns and an additional two or more weeks. While the anomaly team did its work, three maneuvers—designated AMF-2B, AMF-3, and AMF-4—were performed using ACS thrusters only. All of these maneuvers were successful, though it should be noted that the attitude error during these maneuvers showed a more pronounced tendency towards a two-sided limit cycle in the Z axis.

Figure 13 shows one of the data reconstructions performed by the anomaly team. This shows an estimate of the propellant slosh torque acting on the spacecraft. It was derived by calculating a system torque from the system momentum change and subtracting an estimate of the torque imparted by the ACS thrusters and the misalignment of the main engine. Oscillatory propellant slosh dynamics can be seen in both the Y and Z axes of the spacecraft, with the largest effect shown in the Z axis. Note from this plot that it appears that the propellant slosh is damping out, indicating that there may not be an inherent instability to the spacecraft DeltaV Mode because of slosh and that AMF-2 may not have aborted if the FDC limits had been able to be safely set a little higher.
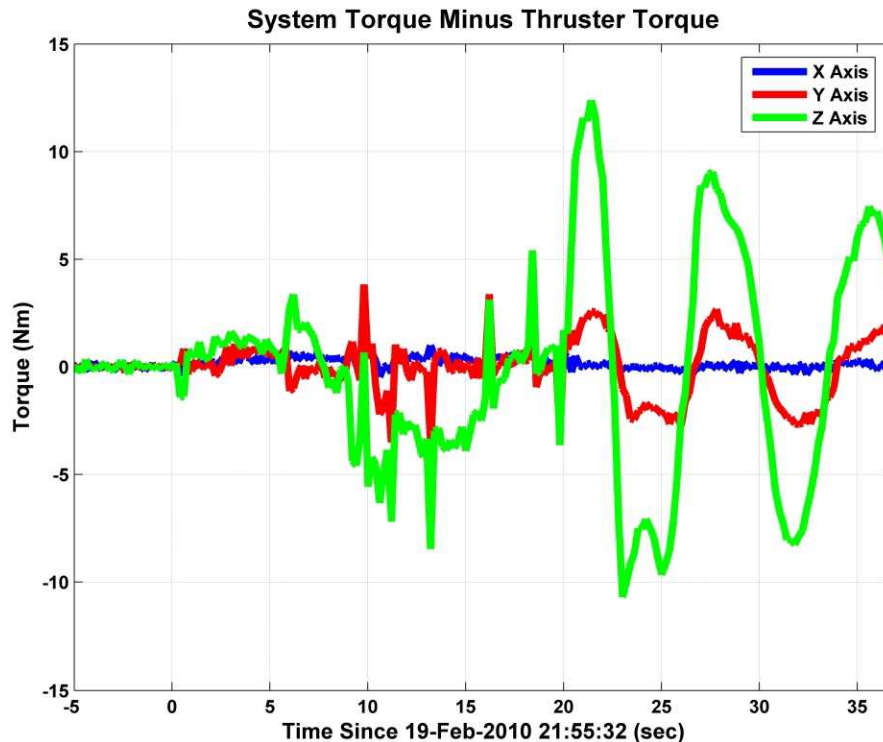

Figure 13: AMF-2 Estimated Propellant slosh Torque

Most of the prelaunch propellant slosh analysis for SDO had been done using a fill fraction of 18–25% and was mainly concerned with the effect of slosh on the jitter seen by the science instruments. The fill fraction at the beginning of AMF-1 was 96% and for AMF-2 was 80%. Based on the pre-launch analysis, we expected propellant slosh effects to increase as the fuel tanks emptied and to reach a maximum around a 60% fill fraction. The anomaly team added a slosh model to the spacecraft high-fidelity simulation (HiFi) and attempted to tune its parameters to match the observed behavior of the spacecraft from the AMF-2 burn (and then the AMF-2B, AMF-3, and AMF-4 burns). Their modeling was not able to completely replicate the observed dynamics, but it did show the correct trends and gave them the confidence to begin developing and testing potential mitigations.

After a lot of time and effort, the anomaly team proposed three simple, easy-to-implement changes to the spacecraft that they felt would solve the problem. First, they suggested increasing the length of the settling burn from 20 seconds to four minutes. That length was selected to encompass the "geyser mode" that can be seen upon initial thruster firing, but did not allow enough time for fuel disturbances to damp out before the main engine firing. Second, the team suggested raising the FDC limit on system momentum to 34 Nms. The 20 Nms limit was chosen based on a two-failure scenario (how much momentum could be placed into the spacecraft by a stuck-on thruster and have the spacecraft successfully recover with one failed wheel); 34 Nms reflected a more realistic failure case. Finally, the team wanted to disable the effect of the structural filter, included in DeltaV Mode to allow it to meet the design requirement for 12 dB modal suppression of all flexible modes, by appropriately setting its parameters to ones and zeros. The increased phase delay caused by the structural filter was affecting DeltaV Mode's ability to react to the slosh dynamics, and the observed flexible modes and damping seen during the appendage deployments (refer back to Figure 4 and Figure 6) indicated that this could be done safely.

AMF-5 was the first post-anomaly maneuver to be performed with the main engine. In order to mitigate the operational effects of another failure, this maneuver was done as a "hybrid" maneuver. The 50 minute maneuver was done as a 40-minute, ACS-thruster-only settling burn followed by the final 10 minutes using the main engine. This was done so that if using the main engine caused the maneuver to fail once again, the overall maneuver would still impart a fair amount of ΔV to the spacecraft. To the relief of everyone on the project team, AMF-5 was a success. In fact, the changes made to mitigate the effect of the slosh dynamics—primarily the removal of the structural filter—significantly improved the response of the controller. After the success of AMF-5, three more AMFs were performed. Figure 14 and Figure 15 show the attitude error, phase plane plots, and system momentum magnitude for AMF-6. As can be seen, the removal of the structural filter from the DeltaV Mode controller significantly improved the response of the spacecraft attitude, resulting in a low magnitude one-sided limit cycle. The phase plane plots (plotted using the same scale as that used for Figure 11 for AMF-1) show much smaller limit cycles for all axes, and lower overall system momentum variations.

A total of 13 maneuvers (including the aborted AMF-2) were performed to move SDO into its mission orbit. After the engineering burn, the successful AMF-1 and the aborted AMF-2, three "AMFs" were performed using ACS thrusters only, followed by the main engine return-to-service hybrid AMF-5, three additional AMFs, and three TMF trims burns. The propellant slosh anomaly resulted in an additional three maneuvers and a delay of only one week in achieving SDO's mission orbit. On March 16, at the conclusion of TMF-3, SDO achieved its final science orbit.

Figure 14: AMF-6 Attitude Error

Figure 15: AMF-6 Phase Plane and System Momentum Magnitude

ATTITUDE CONTROL SYSTEM COMMISSIONING

While SDO's orbit raising activities were ongoing, and for another six weeks afterward, the commissioning, testing, and calibration of the ACS sensors, actuators, and control modes continued in parallel. The bulk of these tests consisted of ACS control mode checkouts and calibration slews used to calibrate the different sensors and actuators, as well as the science instruments and the AIA guide telescopes. In general, the results of these tests showed good performance of the ACS system, meeting all of its requirements. In the remainder of this section, we will discuss some of the interesting things discovered during the commissioning process, including some small performance puzzles that were solved.

Gyro Bias Oscillations

The day after launch, the major activities were the checkout of Inertial Mode, followed by a series of slews using that mode to provide data to calibrate the IRUs. Both of those activities were successful, but small spacecraft oscillations were observed in the IRU bias estimates calculated by the onboard Kalman filter, as shown in the plots on the left hand side of Figure 16. Because the oscillation was at such a low frequency, with a roughly 30 minute period, they were well within the bandwidth of the controller, so the controller followed the oscillating error in the attitude estimate, resulting in actual spacecraft motion. The fact that the spacecraft was actually tracking these oscillations is shown in the guide telescope error signals on the left-hand side of Figure 17. Data from these tests, from subsequent tests of SDO Science Mode, and from all other suitable activities were examined. The oscillations seen on all of the different tests were approximately 10–20 arcsec and would not cause SDO's Science Mode to violate its pointing requirement of [35,70,70] arcsec, but they remained a concern to the science instrument teams. It was determined by looking at this data that the IRUs were the cause, and that the frequency and the amplitude of the oscillations was temperature related.

Because of concerns of the effect of the 87 Hz frequency of the IRU's internal heaters on the spacecraft power bus, the decision had been made not to use them in flight except during a few crucial operations. As a result the IRUs experienced temperature variations that were different from the manufacturer's design and testing profile. That these temperature variations caused the bias oscillations was confirmed by first running a test with the IRU internal heaters turned on; the results of this are shown in the bias estimates and guide telescope signals in the plots on the right hand sides of Figure 16 and Figure 17. Note that because of different average temperatures, the secular component of both the gyro bias estimate and guide telescope error signals are different in the unheated and heated cases, but the scales of each side-by-side plot are the same to show the lower magnitude of the oscillatory component. This test established that the bias oscillation could be reduced through use of the IRU internal heaters, but concerns about the effects on the spacecraft power system remained.

After the first IRU internal heater test, the flight support team ran two tests instead using software-based control of the IRU heaters, first with a setpoint of 40 C and then with a setpoint of 67 C. The software control was relatively coarse, allowing peak-to-peak variations about the setpoint of roughly 2 C. As expected, changing the temperature setpoints resulted in changes to the amplitude and frequency of the observed oscillations. However, they did not significantly decrease the magnitude of the oscillations.

Following the heater tests and the examination of their results, the team decided that the lowest-impact solution to the problem, if possible, would be to adjust the gains of the Kalman filter to make it less sensitive to the low frequency oscillations of the gyro biases caused by the thermal variations. After much analysis and simulation, a new set of gains was available for test.
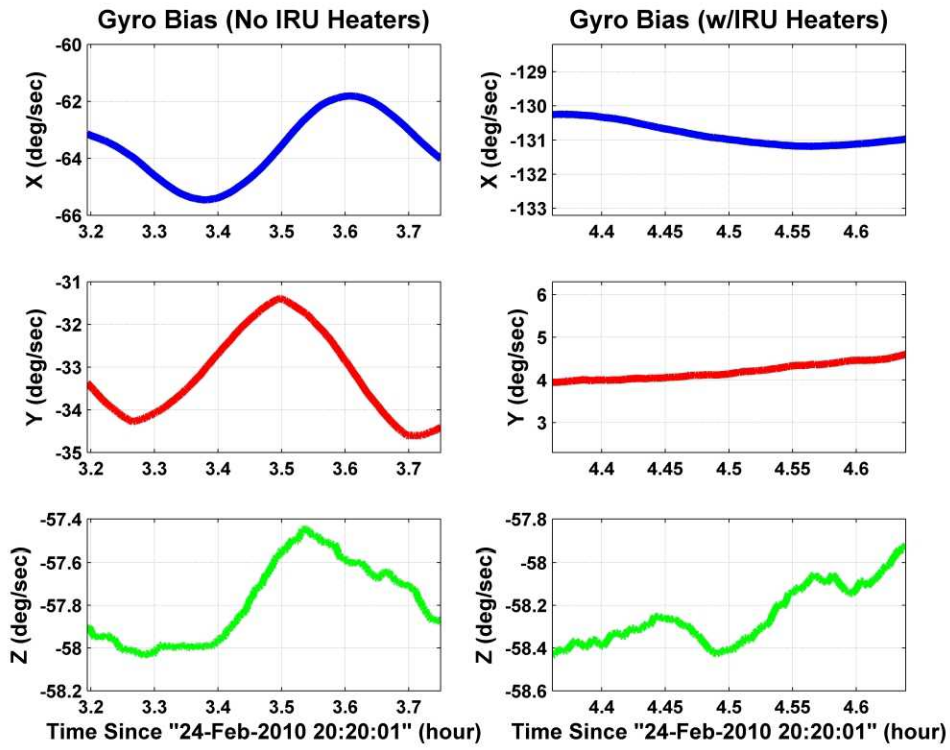
Figure 16: Gyro Biases with and without IRU Heaters


Figure 17: Guide Telescope Signals with and without IRU Heaters

Figure 18 shows a plot of the estimated gyro biases before and after the adjustment of the Kalman filter gains. As can be seen, while the retuned filter does not eliminate the oscillations completely—certainly not as much as using the IRU internal heater—it does diminish their magnitude significantly, bringing the actual oscillation of the Observatory down to a level acceptable to the SDO science teams.


Figure 18: Gyro Biases before and after Adjusting Kalman Filter Gains

Coarse Sun Sensor Sun Angle Calculation

As discussed in the previous section, during the engineering burn and AMF-1, the FDC limit for a Sun vector mismatch (a three-way comparison between CSSA, CSSB, and AD calculated vectors) tripped multiple times. Though the persistence of these limit violations were not nearly long enough to trip any autonomous FDC actions, the team was interested in understanding why they were occurring.

Figure 19 shows the unit sphere of possible Sun vectors in the SDO body reference frame. The dark lines superimposed on the sphere represent the attitude path taken by SDO as it slewed out to the burn attitude for AMF-1 and then back to pointing at the Sun afterwards. The colors of the points on the sphere represent how many individual CSSs register sunlight at each point. The green areas are regions of 4-sensor coverage, the blue areas are regions of 3-sensor coverage, and the cyan areas are regions of 2-sensor coverage; different algorithms are used to calculate the Sun vector for the 4-sensor case than the 3- and 2-sensor cases. During nominal science data collection, four CSSs have visibility of the Sun, and it is in this condition where the calculation of the Sun angle is the most accurate. As shown on the figure, there are several locations along the path where only two or three CSSs can see the Sun. It is in these zones that the limit trips occurred. On both ACEs, the times of CSSA vs CSSB mismatches occurred during times of decreased Sun

20

coverage. In the regions where fewer than four sensors are used, the comparison between CSSA and CSSB is more likely to be different, primarily because the two ACEs may transition between the 4-sensor and 3-sensor algorithms at a slightly different times.


Figure 19: CSS Coverage Map During AMF-1 Slews

ON-ORBIT JITTER TESTING

Two of the SDO instruments, Atmospheric Imaging Assembly (AIA) and Helioseismic and Magnetic Imager (HMI), are sensitive to high frequency pointing perturbations and have sub-arcsecond level line-of-sight (LOS) jitter requirements. Extensive modeling and analysis efforts were directed in estimating the amount of jitter disturbing the science instruments.[7] A jitter test plan was developed to verify analysis approaches and reduce uncertainties in the models. This plan consisted of hardware component tests, structural component tests, and assembly level tests. The hardware component test results were used for calibrating the disturbance input models that drove the observatory structural and optical models used to predict the LOS jitter performance.

Although detailed analysis and assembly level tests were performed to obtain good jitter predictions, there were still several sources of uncertainty in the system. The structural finite element

21

model did not have all the modes correlated to test data at high frequencies (>50 Hz). The performance of the instrument stabilization system was not known exactly, but was expected to be close to the analytical model. A true disturbance-to-LOS observatory level test was not available due to the tight schedule of the flight spacecraft, the cost in time and manpower, difficulties in creating gravity negation systems, and risks of damaging flight hardware. To protect the observatory jitter performance against model uncertainties, the SDO jitter team devised several on-orbit jitter reduction plans in addition to reserving margins on analysis results. Since some of these plans severely restricted the capabilities of several spacecraft components (e.g. wheels and high gain antennas), the SDO team performed on-orbit jitter tests to determine which jitter reduction plans, if any, were necessary to satisfy science LOS jitter requirements.

Post-SDO launch, a three-day test period was dedicated to measure AIA and HMI LOS jitter induced by various spacecraft and instrument mechanisms. In the following sections, brief descriptions of the reaction wheel and high gain antenna jitter tests and samples of on-orbit measurements are provided. More detailed jitter testing information, pre-flight prediction versus measurements, and instrument mechanism test data will be subjects of a separate paper.

Reaction Wheel Jitter Test

Based on pre-flight jitter analysis estimates, the wheel speeds were limited to ±400 rev/min (RPM) to meet AIA and HMI jitter requirements.[8] This constraint forces the wheels to reverse direction about once every four weeks and uses only a small portion of the available wheel speed range of ±6000 RPM. As the wheel reverses direction, a small increase in LOS motion would occur and was considered an acceptable error to the science team. However, the team preferred to reduce the frequency of wheel reversals and the number of thruster maneuvers required to unload wheel momentum. The objective of the wheel jitter tests was to measure the actual wheel-induced jitter on orbit, which was expected to be smaller than the analytical prediction.

During the wheel jitter test, all wheel speeds were driven from a low speed to a maximum speed of about 1000 RPM, and then driven back down to low speed levels. Wheel speed changes were achieved by commanding wheel torques along the null space vector direction to avoid interference with the attitude control torque commands. The primary jitter sensor for this test was the HMI limb sensor measuring LOS errors at 512 Hz. The four AIA guide telescope signals also provided data with sample rates of 256 Hz for telescopes #1 and #3, and 128 Hz for telescopes #2 and #4. Both HMI and AIA instruments can store 3 minutes of data and require 30 minutes to download the data to the ground. To accommodate the instrument data storage and download capabilities, the wheel speeds were accelerated during data collection periods (~3 min) and held constant during data dumping periods (~30 min). Furthermore, during the speed acceleration periods, the acceleration level was small enough (~1 RPM/sec) to ensure that structural modes would be fully excited (e.g. reach maximum amplitude), similar to nominal operating conditions where the wheel speeds change slowly.

The HMI instrument stabilization control system (ISS) was opened while wheel speeds increased to ~1000 RPM and closed while wheel speeds decreased back to low levels. Since HMI nominally functions with a closed-loop instrument control system, jitter data collected during the speed decrement section was compared to the requirement. In Figure 20, the top plot shows the largest AIA LOS measurements from all four telescopes, the middle plot shows the HMI LOS measurement, and the bottom plot shows the corresponding absolute wheel speed levels for four wheels. In the plot legend, LOS Y and LOS Z refer to the tip and tilt motion of the instrument LOS vector.
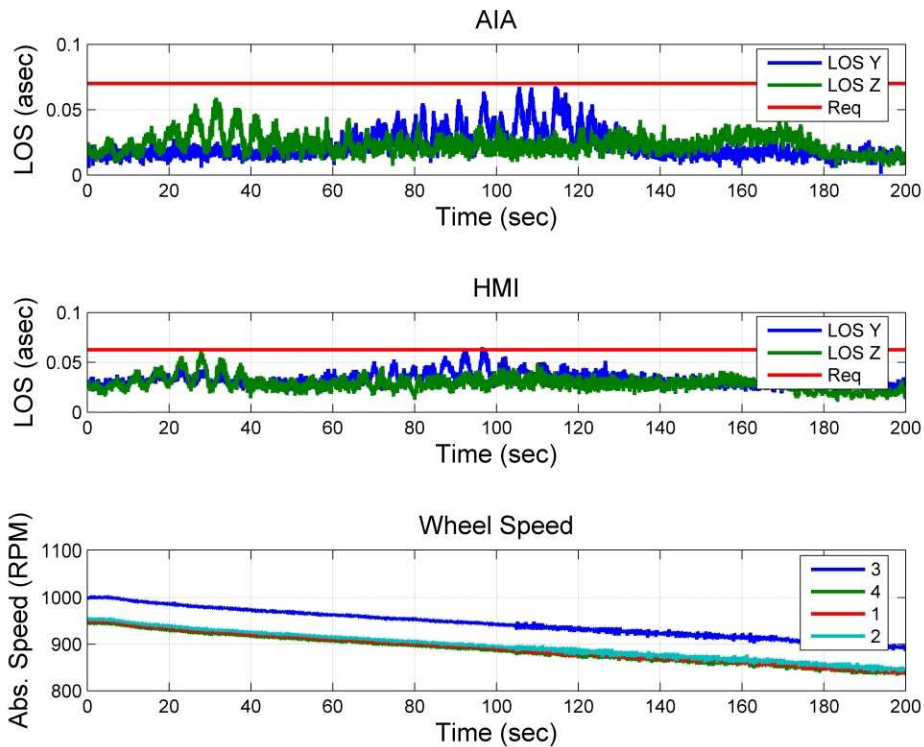
Figure 20: Wheel Jitter Results: AIA/HMI LOS Measurements and Wheel Speeds

These jitter measurements demonstrated that the AIA LOS requirement was met for all wheel speeds tested, whereas the HMI LOS requirement was violated when the lowest wheel speed reached ~840 RPM. Based on these results, the SDO team decided to extend the wheel speed limit from ±400 RPM to ±800 RPM, and thereby reduce the number of direction reversals each wheel will experience over its lifetime and prolong the period between thruster momentum unloads.

High Gain Antenna (HGA) Jitter Test

The goal of the HGA tests was to characterize the jitter induced by HGA operation, and if possible, remove constraints on HGA operation if induced jitter is sufficiently low. The objective was to operate the HGA actuators in all four gimbal axes for a range of representative pulse rates to determine the conservatism, if any, in the pre-flight jitter predictions.

Pre-flight jitter analysis demonstrated that the peak jitter occurred when steps from multiple actuators interfered constructively.[7] In rare cases, this interference resulted in a doubling of induced jitter. To avoid this phenomenon, a stagger-step algorithm was implemented to prevent actuators on the +Z and -Z antennas from taking steps during the same 200-ms HGA control cycle.[9] Even with the stagger-step implementation, high jitter was still predicted to occur very infrequently (less than 1% of the time) and for short durations. AIA was relatively insensitive to this effect, since the resulting image loss could fit within the imaging data continuity budget. However, HMI has a registration requirement that all of the images in a sequence must be taken in a suitable jitter environment. Therefore, the peak jitter was not acceptable. The operational constraint imposed in order to meet HMI requirements took the form of a No Step Request (NSR) flag, sent from HMI to the ACS system, which requests the ACS to not move any of the HGAS gimbals after a specified period for a specified duration.

The HGAS test sequence had two major components:

- Single-step commands were sent to each of the four actuators, in one of several antenna joint configurations (e.g. various azimuth (AZ) and elevation (EL) gimbal angles), starting from a quiet state. The antennas were positioned and held fixed for several seconds to allow any transient motions to die out. The objective of these tests was to see the peak jitter from a single actuator, which provided a floor (minimum jitter level), as well as providing an estimate for peak jitter (twice the single-step jitter).

- Variable step rate commands were sent to each of the four actuators, within the nominal operating range of 0 to 140 deg/hr (maximum allowable jitter rate). The peak jitter occurs when spacecraft modes above 50 Hz are excited. Therefore, the harmonics of the actuator pulses had to align to a fraction of 0.02 seconds with the jitter-sensitive modes (identified from the single-step response above) in order to reach peak jitter. By increasing the individual gimbal rate from 0 to 140 deg/hr with intermediate rates chosen to match typical tracking rates, the goal was to capture peak HGA-induced jitter during nominal operations.

The worst case AIA and HMI LOS jitter measurements are illustrated in Figure 21 and Figure 22, respectively. In these figures, the first four visible jitter peaks are induced by single gimbal steps executed 10 sec apart, with the AZ gimbal stepping twice and then followed by EL gimbal stepping twice. The variable-step rate test results are shown approximately 15 sec after the last EL gimbal step. For the results shown, a single EL gimbal step violates the HGA-induced jitter requirement, and the LOS jitter gets worse as the gimbal step rate increases.


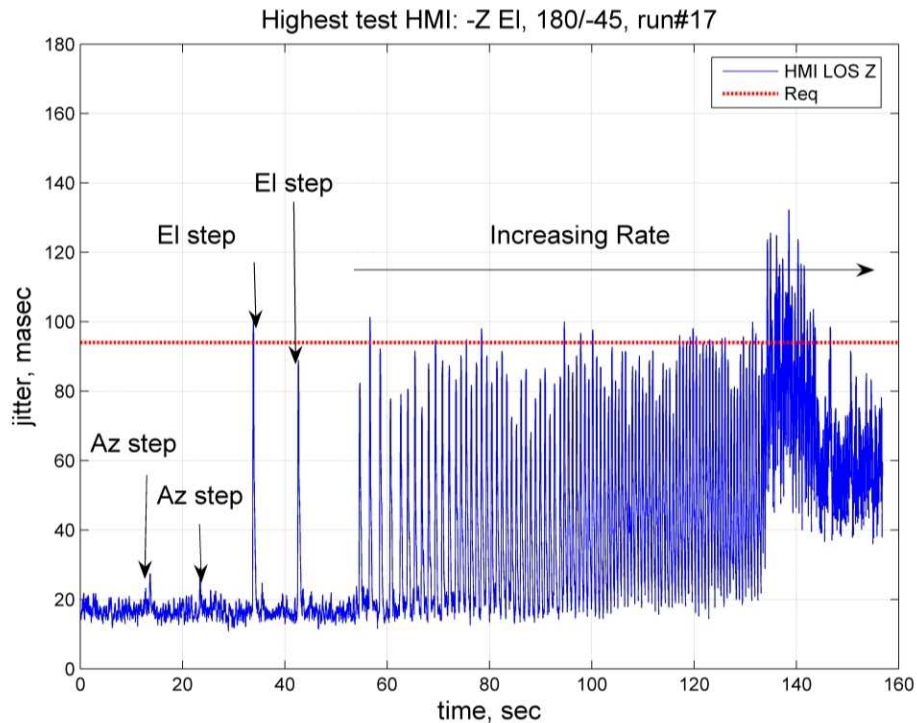**Figure 21: AIA Telescope #3 LOS Measurements from HGA Jitter Test**

Figure 22: HMI LOS Measurements from HGA Jitter Test

Contrary to the team's expectations, the HGA jitter tests showed flight measurements were higher than pre-flight predictions at various HGA joint configurations, despite several levels of conservatism on assumed damping, and the large margins used in the jitter analysis.[10] As a result, HGA jitter mitigation options (stagger-stepping and NSR flags) were enabled to meet instrument LOS jitter requirements. The mitigation algorithms did work as expected and greatly reduced HGA-induced jitter in science images.

CONCLUSION

On May 17, 2010, the SDO Project Manager Liz Citrin "handed over the keys" of SDO to Space Science Mission Operations, signifying the successful conclusion of commissioning and the formal beginning of its science mission. As early as March 31, though, when AIA captured the First Light image shown in Figure 23, SDO was destined for success. Due to the robust design of the ACS subsystem and with the skill and dedication of the flight support team in diagnosing the anomalies seen in flight, the SDO ACS was a key factor in this success.

REFERENCES

[1] Scott R. Starin, Kristin L. Bourkland, Kuo-Chia Liu, Paul A. C. Mason, Melissa F. Vess, Stephen F. Andrews, and Wendy M. Morgenstern. "Attitude Control System Design for the Solar Dynamics Observatory." Flight Mechanics Symposium, 2005.

[2] Scott R. Starin, Melissa F. Vess, Thomas M. Kenney, Manuel D. Maldonado, and Wendy M. Morgenstern. "Fault Detection and Correction for the Solar Dynamics Observatory Attitude Control System." AAS Guidance and Control Conference, Breckenridge, CO, 2008.

[3] Kristin L. Bourkland, Scott R. Starin, and David J. Mangus. "The Use of a Gyroless Wheel-Tach Controller in SDO Safehold Mode." Flight Mechanics Symposium, 2005.

[4] Melissa F. Vess, Scott R. Starin, and Wendy M. Morgenstern. "Use of the SDO Pointing Controllers for Instrument Calibration Maneuvers." Flight Mechanics Symposium, 2005.

[5] Paul A. C. Mason and Scott R. Starin. "SDO Delta H Mode Design and Analysis." International Symposium on Space Flight Dynamics, 2007.

[6] Neil Ottenstein, Greg Natanson, Richard McIntosh, Joe Hashmall, Seth Shulman, Robert DeFazio, Scott Starin, Kristin Bourkland, Paul Mason, and Melissa Vess. "Solar Dynamics Observatory (SDO) Ascent Planning and Momentum Management." Space Ops 2010, Huntsville, AL, 2010.

[7] Kuo-Chia (Alice) Liu, Thomas Kenney, Peiman Maghami, Pete Mule, Carl Blaurock, and William B. Haile. "Jitter Test Program and On-Orbit Mitigation Strategies for Solar Dynamics Observatory." International Symposium on Space Flight Dynamics, 2007.

[8] Kuo-Chia Liu, Peiman Maghami, and Carl Blaurock. "Reaction Wheel Disturbance Modeling, Jitter Analysis, and Validation Tests for Solar Dynamics Observatory." AIAA Guidance, Navigation and Control Conference and Exhibit, August 2008, Honolulu, Hawaii. AIAA 2008-7232.

[9] Kristin L. Bourkland, Kuo-Chia (Alice) Liu, and Carl Blaurock. "A Jitter-Mitigating High Gain Antenna Pointing Algorithm for the Solar Dynamics Observatory." International Symposium on Space Flight Dynamics, 2007.

[10] Carl Blaurock, Kuo-Chia Liu, and Peter Mule. "Solar Dynamics Observatory (SDO) HGAS Induced Jitter." 49th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2008, AIAA 2008-1951.

Figure 23: First Light Image from the SDO AIA Instrument taken on March 30, 2010