

# Investigating On-Orbit Attitude Determination Anomalies for the Solar Dynamics Observatory Mission

Melissa F. Vess\*  
and Scott R. Starin †  
and Dr. Kuo-Chia (Alice) Liu ‡

*NASA Goddard Space Flight Center, Greenbelt, MD, 20771, USA*

The Solar Dynamics Observatory (SDO) was launched on February 11, 2010 from Cape Canaveral on an Atlas V launch vehicle into a geosynchronous transfer orbit. SDO carries a suite of three scientific instruments, whose observations are intended to promote a more complete understanding of the Sun and its effects on the Earth's environment. After a successful launch, separation, and initial Sun acquisition, the launch and flight operations teams began a commissioning campaign that included, among other things, checkout and calibration of the fine attitude sensors and checkout of the Kalman filter (KF) and the spacecraft's inertial pointing and science control modes. During KF and controller checkout, several interesting observations were noticed and investigated.

Initial checkout of the fine attitude sensors indicated that all sensors appeared to be functioning properly. Initial calibration maneuvers were planned and executed to update scale factors, drift rate biases, and alignments of the inertial reference units (IRUs). After updating the IRU parameters, an oscillation was noticed in the sensor residuals and the KF estimation of the IRU bias. A concentrated investigation ensued to determine the cause of the oscillations, their effect on mission requirements, and how to mitigate them. Results of that investigation are presented in the paper.

While the gyro oscillation investigation was ongoing, another unexpected behavior was noticed in one of the star trackers. The ST2 quality index displayed an unusual shape. Instead of remaining fairly constant around the expected value of 0.95 (in a range of zero to one), there was a period of approximately six hours during which the quality index showed an inverted "U" shape, dropping sharply to as low as 0.6 at the beginning of the period, climbing back towards 0.95 during the middle of the period, and then dropping again, similar to the start of the phenomenon. The poor quality solutions did not adversely affect the on-board spacecraft attitude estimate because the flight software already included logic that rejects any ST solutions whose quality index is below a database threshold. Still, the very distinctive shape of the quality index was curious enough to warrant further investigation, carried out jointly between the Attitude Control Systems team and the Attitude Ground Systems team. The methods of collecting the necessary data for investigation and preliminary results of the investigation are presented in the paper.

## I. Introduction

The Solar Dynamics Observatory (SDO), depicted in Figure 1, launched on February 11, 2010 from Cape Canaveral on an Atlas V launch vehicle into a geosynchronous transfer orbit. SDO carries a suite of three scientific instruments, whose observations are intended to promote a more complete understanding of the Sun and its effects on the Earth's environment. The Helioseismic & Magnetic Imager (HMI) is designed to study oscillations and the magnetic field at the Sun's surface. The Extreme Ultraviolet Variability Experiment (EVE) measures the solar extreme ultraviolet irradiance in order to understand how variations in the timescales influence Earth's climate. The Atmospheric Imaging Assembly (AIA) consists of four science telescopes that image the Sun in various wavelengths in an effort to understand the physics behind the solar surface activity.<sup>1</sup> In addition to their science telescopes, AIA also has four fine Sun sensors (guide telescopes) that are used both to drive the image stabilization systems (ISS) on the science telescopes, and as an attitude sensor in the spacecraft's on-board flight software (FSW). One of the guide telescopes, designated the controlling guide telescope (CGT) is used as the primary pitch/yaw sensor in the Science pointing attitude control mode.

After a successful launch, separation, and initial Sun acquisition, the launch and flight operations teams dove into a commissioning campaign that included, among other things, a series of orbit raising maneuvers to bring the observatory to its final geosynchronous orbit, checkout and calibration of the fine attitude sensors, checkout of the Kalman Filter (KF) used for on-board attitude determination, checkout of the spacecraft's Inertial pointing and Science

\*Aerospace Engineer, Attitude Control Systems Engineering Branch, Code 591, AIAA Member.

†Aerospace Engineer, Attitude Control Systems Engineering Branch, Code 591, AIAA Member.

‡Aerospace Engineer, Attitude Control Systems Engineering Branch, Code 591, AIAA Member.

control modes, and initial calibration of the Science instruments.<sup>2</sup> Early on in the commissioning campaign, high spacecraft momentum resulted in an abort of one of the orbit raising burns. For the next two weeks, much of the attitude control system (ACS) team's focus shifted to troubleshooting what turned out to be a fuel slosh incident.<sup>3</sup> Concurrent with the fuel slosh investigation, the remaining ACS team members continued to support checkout and calibration of the fine attitude sensors and on-board Kalman Filter, as well as checkout of the Inertial and Science mode controllers. During that process, several interesting observations were noticed and investigated.

The SDO fine attitude sensors consist of one Adcole Digital Sun Sensor (DSS), two Galileo Avionica quaternion-output Star Trackers (STs), and three Kearfott Two-Axis Rate Assemblies (inertial reference units, or IRUs). Initial checkout of the fine attitude sensors indicated that all sensors appeared to be functioning properly. Initial calibration maneuvers were planned and executed to update scale factors, biases, and alignments of the IRUs. After updating the IRU parameters, the Kalman Filter was initialized and quickly reached convergence. Over the next few hours, it became apparent that there was an oscillation in the sensor residuals as well as the Kalman Filter estimation of the IRU bias. A concentrated debugging effort ensued to determine the cause of the oscillations and how to mitigate them. Section III. describes the investigation process and subsequent results and mitigations.

While that investigation was on-going, another unexpected behavior was noticed in one of the star trackers. The ST2 quality index, a scalar value that represents the weighted average of the axis-by-axis solution covariance, displayed an unusual, inverted "U" shape. The quality index can range in value from 0 to 1, where 1 represents an almost perfect solution. Instead of remaining fairly constant around a value of 0.95 (the nominal value seen up to that point), over a 6 hour period, the quality index dropped suddenly to a value as low as 0.6, slowly increased again to 0.95, and then slowly decreased to 0.6 again. Section IV. describes the investigation into the low quality index that ensued.




Figure 1. Solar Dynamics Observatory

## II. Background

### A. Inertial and Science Mode Pointing

The SDO Inertial and Science mode controllers are designed predominantly to keep the spacecraft fixed on the center of the Sun. As mentioned before, the Science mode controller uses one of AIA's guide telescopes to directly control the pitch (Y) and yaw (Z) axes of the spacecraft. For the Science mode roll (X) axis and Inertial mode, however, control is accomplished by calculating the error between the estimated spacecraft attitude (discussed in Section B. below) and the very slowly time-varying Solar North Reference (SNR) frame target attitude. The SNR frame aligns the spacecraft X axis with the center of the Sun, aligns the spacecraft Y axis with the cross product of the X axis and the Solar North Pole, and aligns the spacecraft Z axis with the cross product of the X axis and the Y axis.<sup>4</sup> The SNR frame rotates at approximately 0.04 asec/sec, mostly about the spacecraft Z axis. The IRU measured rates are used for both the Science and Inertial mode controllers to damp spacecraft rates. When the spacecraft is fixed on its target, the measured rates are essentially zero.

### B. Attitude Determination

SDO has an attitude determination requirement of [75, 90 90] asec,  $3\sigma$ , of which, [35, 70, 70] asec,  $3\sigma$  is the ACS allocation. The tight knowledge requirement is needed in order to ensure that the error between the Inertial mode pointing, which uses an estimate of the spacecraft attitude, and the Science mode pointing, which is based directly off of the CGT, is small enough to allow for the spacecraft to point within the small (less than 95 asec) field of view of the CGT while in Inertial mode. The close co-alignment between Inertial and Science modes is necessary for several reasons. First, it allows for a smooth transition from Inertial to Science mode as the spacecraft moves from not using the CGT to pointing based off of the CGT measurements. Second, most of the instrument calibration maneuvers are performed in Inertial mode, so the desire is to have the attitude with respect to the Sun very similar to what the instruments see in the Science pointing mode. Finally, SDO has regular eclipses during periods of the year. While in eclipse, the spacecraft cannot remain in the Science pointing mode because the CGT measurements will become invalid. The spacecraft will have to enter Inertial mode prior to the eclipses, and the Science teams want to stay as close to nominal Science pointing as possible prior to and during eclipse to minimize the time during which Science data cannot be gathered.

In Inertial mode and for the Science mode X axis, the spacecraft attitude is estimated using a six-state, extended Kalman Filter. The KF estimates adjustments to the spacecraft attitude and IRU biases.<sup>4</sup> The inputs to the KF algorithm are the IRU-measured rates, ST-measured attitude quaternions, and DSS-measured unit Sun vector. For redundancy, SDO has three two-axis gyros, but at any given time, the FSW is configured to use three of the six IRU axes. The KF estimates biases only for the three primary IRU axes. At any given time, only two of the three fine attitude sensors are needed to meet the attitude knowledge requirement.

### C. SDO Inertial Reference Units

As mentioned above, the SDO IRUs are three Kearfott, two-axis rate assemblies. The Kearfott IRUs have an internal temperature controller, designed to maintain the IRU temperature at a constant temperature of approximately 70°C. The heater control is designed to minimize the change in the bias drift and scale factors of the mechanical gyros. The IRU analog heater control turns the heaters on and off at approximately 83 Hz. During ground testing of the observatory, it was discovered that the 83 Hz control cycle of the IRU heaters disturbed the Power System Electronics (PSE) battery current controller. With a bandwidth at around 100 Hz, the charge/discharge controller tracked the power bus oscillation caused by the IRU heater switching, varying the battery charge current by 1.5 A. The SDO flight batteries were never qualified to operate with this type of variation in the fine control of the charge/discharge loop, so, ultimately, the decision was made to baseline leaving the IRU heaters powered off rather than risk battery life and reliability.

### D. SDO Star Trackers

The SDO STs are quaternion-output star trackers that can acquire their attitude, based off of star measurements, from anywhere in the celestial sphere. The 16.4° by 16.4° field of view (FOV) of the each ST is made up of charge couple devices (CCDs) consisting of 512 × 512 pixels. The STs are mounted on the Observatory instrument module and are aligned such that the boresights are not coaligned. Each boresight is rotated approximately 26 degrees from the +Z spacecraft axis, and thus the two boresights are approximately 52° apart. In this orientation and when the spacecraft is pointed at its SNR target, each ST will go through a period of six months during which it will have daily occultations by the Earth. The length of the occultations will start small and will increase to approximately four hours before decreasing again.

The STs have seven software modes: Initialization (INI), software maintenance (SWM), standby (STB), autonomous attitude determination (AAD), single shot attitude (SSA), CCD image (CDI), and commanded star tracking (CST). After power on, the ST starts in INI mode. It can then be commanded into STB mode, which starts running the health and safety checks on the ST. These health and safety check are performed continuously whenever the ST is in STB or AAD mode. SSA is a debugging mode in which the ST takes a picture of the FOV and provides for downlink all of the stars (up to 48) seen in the ST FOV. CDI mode is another debugging mode, in which the ST can be commanded to measure the energy level of each pixel within a small "window" (up to 3056 pixels at a time) of the CCD. CST is a third debugging mode that will track specific stars designated in the command to CST mode. SWM mode is used for any necessary maintenance (upload or download) of the internal software.

AAD mode is where the ST will spend the majority of its time. Upon command, the ST can enter AAD mode from STB mode. If AAD mode is entered from a "lost in space" condition, the ST will scan the FOV to identify stars and triads of stars that can be used to determine the absolute attitude of the ST with respect to the J2000 frame. Once the attitude is determined and verified, the ST will remain in AAD mode and continue tracking up to 9 stars. The star vectors used by the ST to determine and verify attitude are made available for downlink telemetry. The ST provides both the measured stars in the ST frame, as well as which catalog (reference) stars the measured stars are mapped to. The reference star vectors are provided in the Geocentric Inertial (GCI) frame. As part of the AAD mode telemetry, the ST also provides a quality index, which is a measure of the quality of the estimated attitude. The quality index is a value between 0 and 1, where values close to 0 indicate a very poor solution and values close to 1 indicate a high quality solution.

## III. IRU Bias Oscillation

### A. Discovery of the Problem

After a successful initial acquisition of the Sun, the first priority of the Operations team was to checkout SDO's Safehold mode (the safety net of the mission), followed by a checkout of the thrusters, which would be used to help bring SDO into its Science orbit. With those activities accomplished, focus could shift to starting the checkout of the

systems that would ultimately be used during the Science collecting phase of the mission. Up first was initial check out and calibration of the on board attitude determination system. The check out and calibration included checking the functionality of the fine attitude sensors themselves, as well as verifying the functionality of the on-board Kalman Filter software. In addition, a series of spacecraft slews were executed and data from the DSS, STs, and IRUs was used by the Attitude Ground Systems (AGS) team to estimate alignments of the sensors, along with scale factor and bias estimates for the IRUs. The new parameters were loaded to the Spacecraft flight software and functionality of the on-board KF was checked again, at which point the ACS analysis and AGS teams noticed oscillations in the estimates of the IRU biases - both with the on-board KF, and the ground-based KF. Figure 2 shows the IRU bias estimates from the on-board KF for the primary IRU axes.




Figure 2. On-board KF estimates of Primary IRU axes

As can be seen in the figure, the X- and Y-axis bias estimates have approximately the same period of about 30 minutes. The Z-axis bias estimate has a much shorter period of about 5 minutes. In addition, the mean bias of all three axes was shifting over time. Further investigation revealed that not only was the estimated bias oscillating, but the sensor residuals were also oscillating with similar frequencies. Figure 3 shows the ST1 sensor residuals for the same time period. The oscillation in the X axis can be seen clearly, but the Y- and Z-axis oscillations are more difficult to see due to the noise in the ST boresight. Recall that the ST boresight is aligned in the spacecraft YZ plane. ST2 and the DSS showed similar oscillations. In addition, Figure 4 shows that the ST measured quaternion was also oscillating, implying that the spacecraft appeared to actually be following the oscillation. Because the Inertial mode controller uses the estimated attitude for control in all three axes, and because the oscillations were at such a low frequency, they were well within the bandwidth of the controller, and therefore the PID controller faithfully followed the oscillating attitude estimate. Since the IRU rate estimate (and thus the bias estimate) is used in the calculation of the sensor residuals, the question arose as to whether the residuals were oscillating because of the estimated bias oscillation, or if the oscillating residuals were driving the bias estimate to oscillate. Further investigation was needed to determine the source of the oscillation.

## B. Evaluation of the Problem

The investigation into the oscillating bias estimate lead the ACS team to determine that there was an oscillation in the estimated attitude, but at that point in time, the origin of the oscillation was still unknown. Because the KF




Figure 3. ST1 residuals

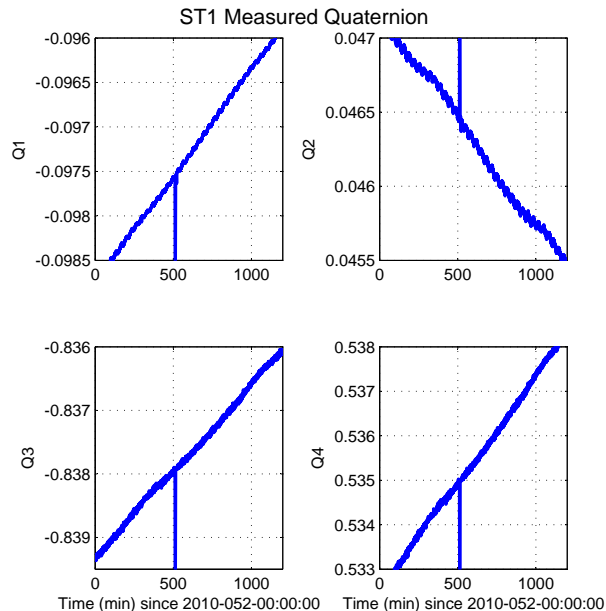


Figure 4. ST Quaternion

uses the IRUs, STs, and DSS to determine an attitude estimate, any one of those sensors could have been driving the oscillation in the attitude and bias estimates. Because the oscillations were being seen in all three axes, it was unlikely that the DSS was the source of the oscillations, which meant the source had to be either the IRUs or one or both of the STs. An attempt was made to isolate those sensors by reconfiguring the Inertial mode controller to control directly off of one of the ST measurement quaternions for attitude error. In doing so, the oscillations in the ST measurement quaternion damped out, implying that the ST was not the source of the oscillation, and thus leaving the IRUs as the culprit. It appeared that the IRU biases really were oscillating, but the Kalman Filter in its current configuration could not keep up with the changing biases. Therefore, the estimated attitude was still oscillating and driving the spacecraft to oscillate as well.

Once the AIA guide telescopes were powered on, the ACS team could actually track the size of the oscillations in the pitch and yaw axes. Figure 5 shows the measurements from one of the AIA guide telescopes. As can be seen, the oscillations during this time frame were on the order of  $\pm 10$  to  $\pm 20$  asec.

The next question the analysis team sought to answer was why the IRU biases were oscillating. Figure 2 showed that the X- and Y-axis bias estimates had the same period of oscillation, while the Z-axis period was different. It also indicated that the mean biases of the three axes were shifting over time. In the current configuration of the spacecraft, the X- and Y-axes both came from IRU-1, while the Z-axis came from IRU-2. Figure 6 shows the internal temperatures of the three IRU boxes. Notice that shifts in the mean values of the bias estimates appear to correspond to shifts in the IRU internal temperatures. Further data analysis confirmed that shifts in the internal IRU temperature resulted in changes to both the amplitude and period of the bias oscillations, though the changes did not appear to be predictable.

As was discussed in Section C., these particular gyros were intended to be run with an internal temperature controller and that concerns over battery charging cycling resulted in a project decision to run the IRUs in an unheated manner. Once it was determined that the characteristics of the bias oscillations were driven by IRU temperature, a test was




Figure 5. Guide Telescope measurements




Figure 6. IRU temperatures

performed to change the IRU temperatures to see the effect on the bias oscillations. As a mitigation to the IRU temperature/battery cycling risk, ACS implemented a mode in which it could use the spacecraft FSW to control the IRU temperatures. The on-orbit test used the FSW heater controller to set the IRU temperatures to 40°C and held them there for about 2.5 hours. The temperatures were then changed to 67°C, which is just below the 70°C set-point of the internal IRU temperature controller, again using the FSW heater controller. The temperatures were held there for just under 2 hours. Finally, the internal IRU temperature controller was used to control the IRU temperatures to the 70°C IRU design temperature. The temperatures were held there for over 3.5 hours. After each change to the IRU temperature, the Kalman Filter was reset and allowed to reconverge on the new IRU biases. The following figures show the IRU temperatures and the estimated IRU biases for each phase of the testing. For the bias plots, the mean bias has been subtracted out each time in order to show more clearly how the amplitude of the oscillations changed as the IRU temperature changed.

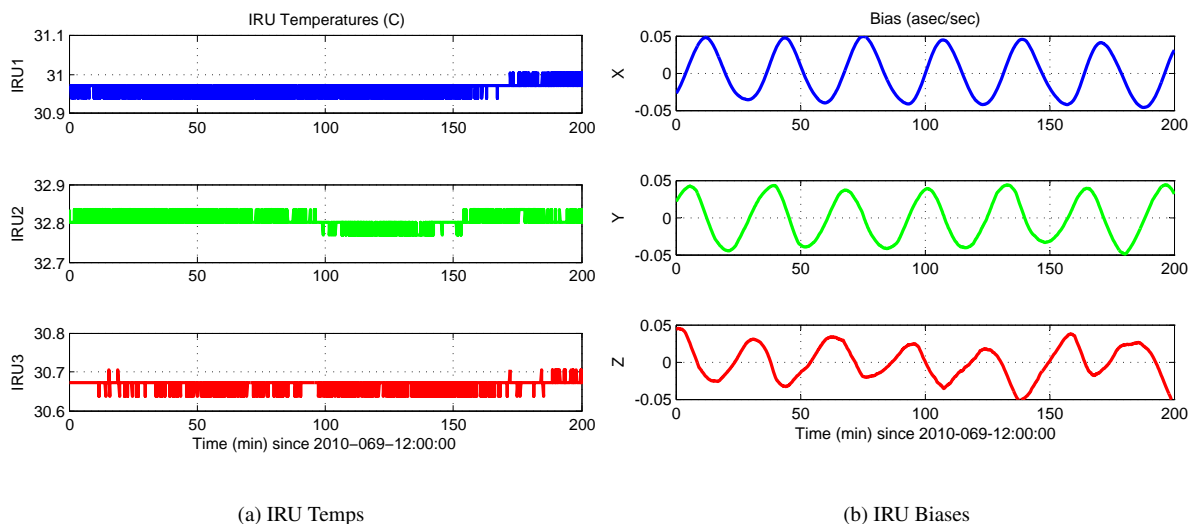


Figure 7. IRU temperatures and KF estimated IRU biases prior to the start of the heater testing

As can be seen in the figures, changing the IRU temperatures changed both the amplitude and the period of the estimated bias oscillation. However, even running with the as-designed hardware heater control did not fully

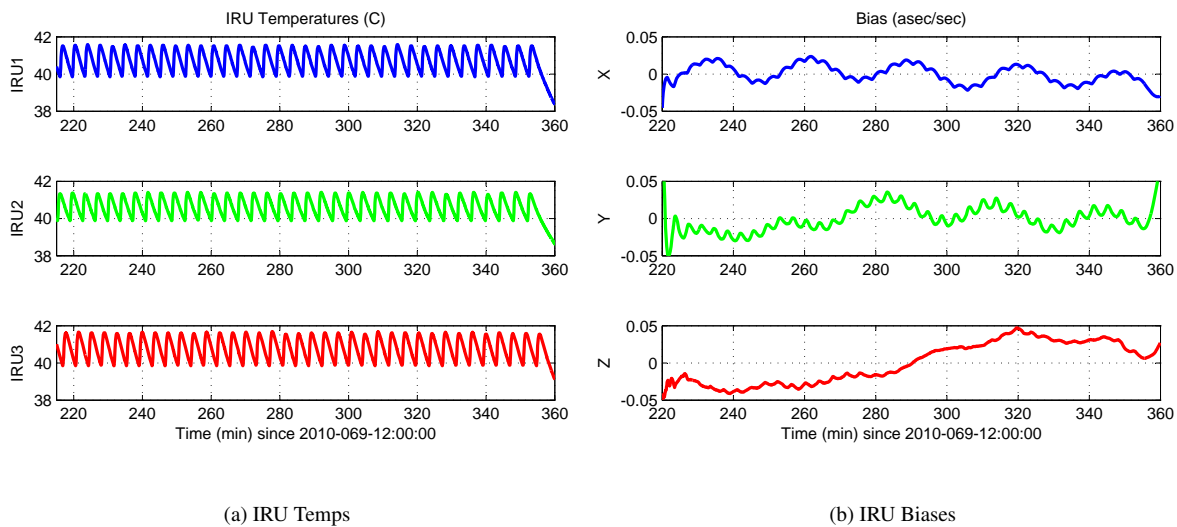


Figure 8. IRU temperatures and KF estimated IRU biases with the software heaters set to 40°C

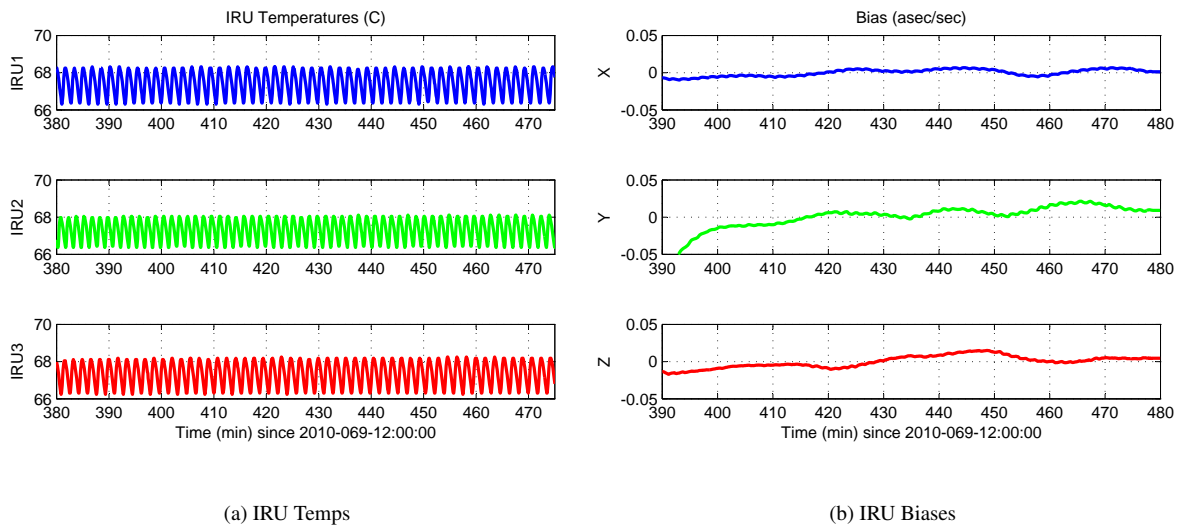


Figure 9. IRU temperatures and KF estimated IRU biases with the software heaters set to 67°C

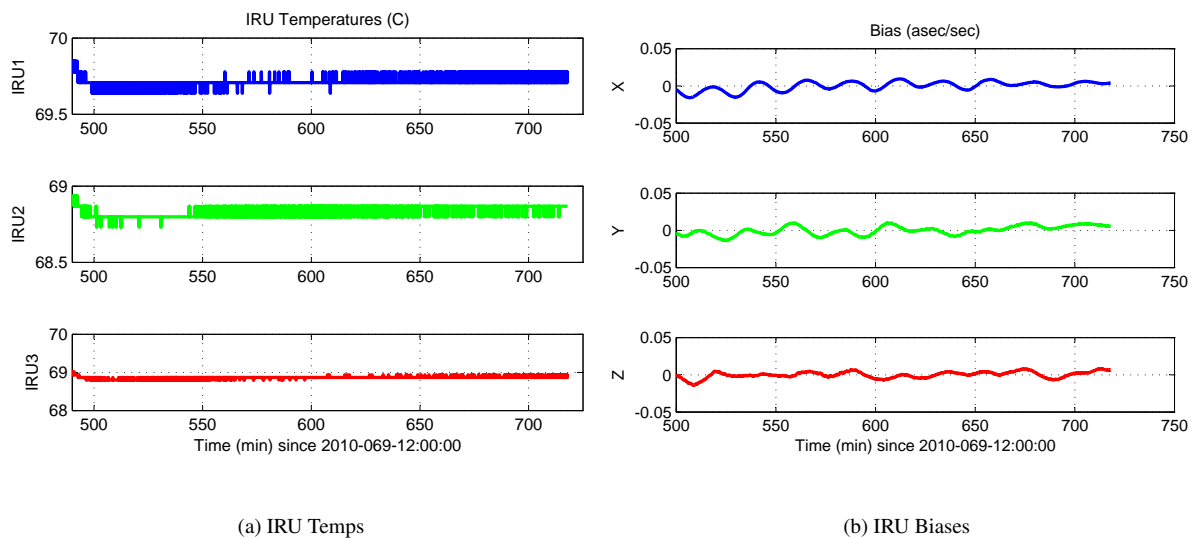


Figure 10. IRU temperatures and KF estimated IRU biases with the hardware heater control enabled

remove the bias oscillation. This fact can be even more clearly seen in Figure 11, which shows the guide telescope measurements for the period of time during which the hardware heater controller was being used. The oscillation in the GT measurements has been reduced, but has not been completely removed.

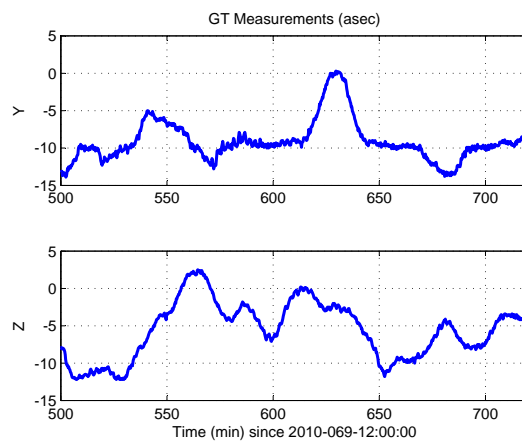


Figure 11. GT Measurements with the IRU hardware heater control enabled

When the bias oscillation first presented itself, the analysis team contacted engineers in the ACS hardware branch as well as the IRU vendor, Kearfott, for their thoughts as to what could be causing the oscillation. One idea presented as a possible cause was a disturbance during launch that caused the gyro flywheel to touch down on a stop. Such a behavior could effect the hysteresis motor, thus disturbing the gyro balance loop. Unfortunately, the data that could be collected on orbit did not provide enough insight into the workings of the gyros, so no further analysis on root cause was possible.

### C. Mitigation of the Problem

Knowing the size of the oscillations and therefore the motion of the spacecraft, the ACS, Operations, Systems, and Instrument teams then discussed the impact of such motion on Science data. The ACS allocation of the attitude knowledge requirement is  $[35, 70, 70]$  asec,  $3\sigma$ . While oscillations on the order of  $\pm 20$  asec were still within



requirements (though just barely in the X axis), the spirit of the requirement assumed that the knowledge error would be a bias with high frequency white noise, or at worst a very long term variation - not an oscillation with a 30 to 40 minute period. When the Science mode controller is used, the effect of the bias oscillation would be minimal for the spacecraft pitch and yaw axes due to the controller error coming directly from the AIA CGT. However, the X-axis oscillations would still be there, and there will be oscillations in all three axes when the Inertial mode controller was used. The Inertial mode controller is used to perform a number of instrument calibration maneuvers. Trying to calibrate the Science instruments while the spacecraft motion was oscillating would be difficult. Based on these considerations, the project teams decided that something needed to be done to mitigate the effects of the bias oscillation.

As was seen in Figures 10 and 11, turning on the IRU hardware heater control did decrease the oscillations in the IRU biases and therefore the motion of the spacecraft. However, it did not completely remove the oscillations. In addition, due to the concern over battery cycling when the IRU heaters are running, the project wanted to pursue other mitigation options. The ACS team came up with several additional options.

The first option considered was to fly directly off of one of the ST quaternions for attitude control. In Science mode, the ST solution would be used only for X- (roll) axis attitude control. The STs are mounted in such a way that they provide their most accurate information about the spacecraft roll axis, so initial analysis indicated that the 35-asec roll axis knowledge requirement could still be met using a single ST solution. In Inertial mode, however, the ST solution would be used for attitude control in all three axes. The ST is very noisy about its boresight, which is mounted mostly along the spacecraft Z-axis. Attitude determination analysis indicated that using a single ST for attitude knowledge would not be sufficient to meet the 70 asec pitch and yaw axis knowledge requirements.

The second additional mitigation option considered was to use the structural filter included within the Science and Inertial mode controllers to notch out the oscillation frequency of the IRU biases and make the controller less responsive to the oscillations. Each controller has a separate structural filter, and, in addition, a different filter can be used for each spacecraft axis. The coefficients for the up-to fourth order filter are FSW parameters, so changing the filter coefficients would mean just a table load to the FSW and would not require a change to the actual FSW code. The major draw back to this option was that it relied on the bias oscillations being fairly constant in frequency. It has already been seen that the bias oscillations can change in amplitude and frequency with changing IRU temperature. Until the spacecraft reached its geosynchronous orbit, the steady state temperatures of the IRUs could continue to change. Even once on station, the daily eclipses seen by the spacecraft during eclipse season would cause the IRU temperature to vary. Long term monitoring of the bias oscillations would be needed to be able to design a notch filter with the correct frequency range. In addition, changing the structural filter for the controllers to accommodate a very low frequency structural filter would necessitate a complete redesign of the controller with extensive stability and performance analysis. Such a change would be time consuming and difficult considering how much time and effort it took to develop the current controller to balance between stability concerns and fast slew and settle requirements. More information on the development of the Science and Inertial mode controllers can be found in Ref. 4.

The final mitigation option considered, and the one subsequently chosen, was to tune the KF to be less responsive to the oscillations in the IRU biases. The ACS High Fidelity Simulator (HiFi), was used to model the oscillating biases and investigate the affects of adjusting different KF parameters. The HiFi is a faster than real time, closed-loop simulator that models the spacecraft dynamics, sensors, actuators, and FSW algorithms, and can be used to analyze spacecraft behavior. The KF parameters modified were the measurement noise (*i.e.* sensor noise) covariances for the two STs and the Digital Sun Sensor (DSS) and the state noise covariance matrix. The state noise covariance matrix was modified by modifying the IRU drift bias parameter,  $\sigma_u$ . The following plots from the HiFi show the estimated attitude error, estimated IRU biases, and true attitude error for various combinations of the sensor noise covariance and state noise covariance. Figure 12 shows the HiFi modeling using the default KF parameters. The period of the oscillation has been adjusted to shorten the run time of the HiFi simulations.

Figure 13 shows what happens to the estimated attitude error, estimated IRU biases, and true attitude error when the sensor noise covariances were reduced by a factor of 10, thus implying that the KF should weight the sensor measurements more heavily than the state estimate. The oscillation in the true attitude error is reduced from that of the default biases, but there still appears to be an oscillation on the order of about 3 asec. The attitude error noise is slightly higher, but not significantly so. Similar results were seen in the HiFi when  $\sigma_u$  was increased by a factor of 10, implying that the KF should weight the state estimate less than the sensor.

Figure 14 shows what happens to the estimated attitude error, estimated IRU biases, and true attitude error when the sensor noise covariances were reduced by a factor of 10 and  $\sigma_u$  was increased by a factor of 10. As can be seen, the oscillations in the true attitude error are virtually eliminated. However, the attitude error noise does increase significantly. The effect of that increased noise on the true attitude error is minimal, however, due to the relatively low bandwidth of the controller. It was ultimately decided to increase  $\sigma_u$  by a factor of 10 and reduce the sensor covariances

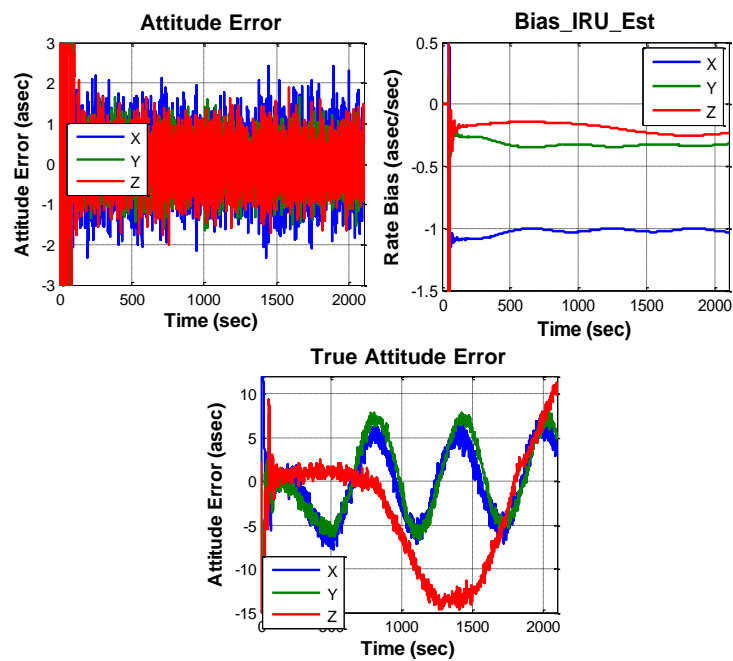


Figure 12. Estimated attitude error, estimated IRU biases, and true attitude errors from the HiFi simulator for the default KF parameters

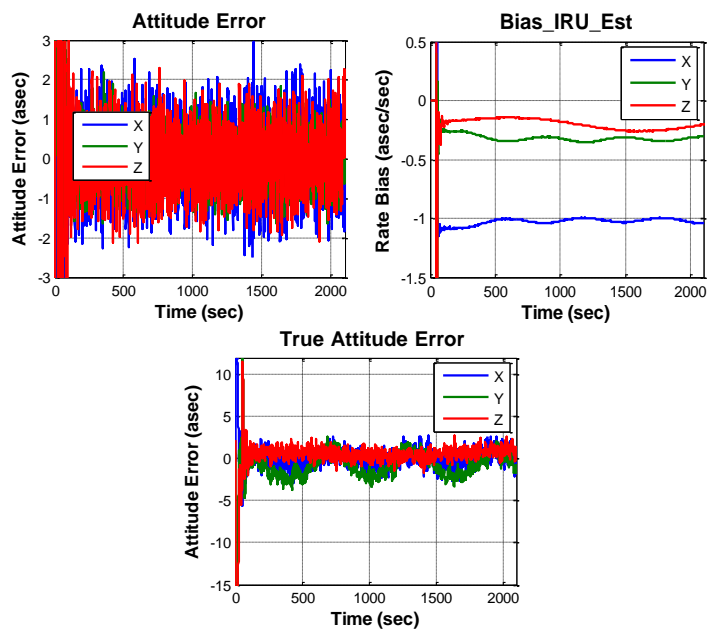
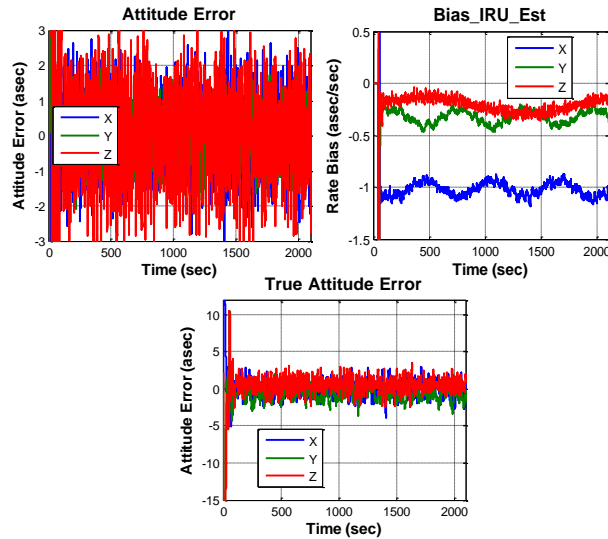



Figure 13. Estimated attitude error, estimated IRU biases, and true attitude errors from the HiFi simulator with the KF sensor residuals reduced by a factor of 10

based on estimates from the AGS team. Figure 15 shows the estimate IRU biases and GT measurements after updating the KF parameters. As can be seen, the oscillation in the GT measurements has been virtually eliminated. The offset from zero is due to the fact that the KF sensors and the GTs had not yet been coaligned.



**Figure 14. Estimated attitude error, estimated IRU biases, and true attitude errors from the HiFi simulator with the KF sensor residuals reduced by a factor of 10 and  $\sigma_u$  increased by a factor of 10.**



**Figure 15. KF estimated IRU biases and GT measurements after updating the KF parameters**

## IV. ST Low Quality Index

### A. Discovery of the Event

On 08 March 2011, approximately one month after the launch of SDO, the ACS and AGS teams noticed that the ST2 quality index was frequently being flagged as “low quality”. Based on feedback from the manufacturer on the relationship between solution quality and reported quality index, the ACS team had already designed a check within the on board FSW to flag ST solutions whose quality index was below a database specified limit (currently set to 0.7) as “low quality”. The nominal performance of the ST up to that time maintained a quality index at or above 0.9 except for times when the Earth was approaching the FOV. The flagging of ST2 solutions as “low quality” was occurring during a time when the Earth was well out of the ST FOV. When the quality index for ST2 was plotted over a period of about 10 hours, an interesting shape of the quality index was noted. Figure 16 shows the quality index for that time frame. Notice that it looks like an inverted “U”.




Figure 16. ST2 quality index with unusual inverted “U” shape.

Coincidentally, a few weeks prior to launch, the SDO team learned of an anomaly on the European Space Agency (ESA) mission, Herschel, whereby their star trackers, also GA star trackers, were accumulating “warm pixels” - pixels whose background energy level were higher than the energy levels of the surrounding pixels.<sup>5</sup> The effect of these “warm pixels” was that when tracked stars were near the warm pixels, it would distort the image of the stars and thus result in a lower quality solution. The signature symptom of a “warm pixel” can be seen by looking at the reported quality index. As a tracked star nears a warm pixel, the quality index will drop due to the higher energy level of the warm pixel blurring the image of the star. As the center of the tracked star moves over top of the warm pixel, the higher energy level of the warm pixel coincides with the center of the star, thereby not corrupting the observed position of the star in the ST FOV, and the quality index improves. Finally, as the tracked star moves away from the warm pixel, the star image again blurs and the quality index drops again. The resulting shape of the quality index over time, as the tracked star moves toward, across, and away from a warm pixel is that of an inverted “U” shape.

### B. Evaluation of the Event

Knowing of the Herschel anomaly, the ACS team had already started working on investigative and mitigation procedures in the event that SDO developed symptoms of “warm pixels”. At the time the SDO event occurred, the cause of the accumulated warm pixels on Herschel was still unknown. For SDO, there was no initial concern over potential “warm pixels” affecting the on board estimated attitude solution. As was briefly described above, the on board FSW already included a mechanism to flag poor quality ST solutions. Such low quality ST solutions are excluded from

use in updating the KF solution. The immediate concern for SDO was in characterizing the problem and in identifying if the inverted “U” quality index was, in fact, caused by a warm pixel.

The main difficulty the SDO team had to overcome was to develop a way to collect the data necessary for trying to determine if the low quality event witnessed was due to a warm pixel. In discussions with the manufacturer, they requested data on the measured attitude quaternion, the stars being tracked and referenced, and the energy levels from the CCD pixels all corresponding to an attitude near which the low quality event was seen. In addition, they also requested the defective pixel table that is maintained internally by the ST FSW. Unfortunately, due to the way the ST telemetry is generated and the SDO FSW 1553 bus polling schedule was developed, all of that data was not readily available in the downlink telemetry. The ST distributes its telemetry across multiple subaddresses of its remote terminal (RT) address on the 1553 bus. SDO’s FSW was set up to poll for the attitude data at 5 Hz, the housekeeping data at 1 Hz, and the measured and reference star data at 1 Hz. The measured attitude quaternion is available within the attitude data, so obtaining that information was not a problem. The rest of the data, however, was problematic to various degrees.

The ST runs on an internal 10 Hz cycle. While in AAD mode, the ST tracks up to 9 stars. The data associated with all 9 stars are packaged within 2 telemetry packets (measured stars and reference stars), 3 stars at a time, each ST cycle. In other words, to downlink data from all 9 stars, the star data subaddresses must be polled 3 times every ST cycle, or at a rate of 30 Hz. With SDO’s current polling schedule, only 1/3 of the star data is downlinked at any given time. An alternative method of obtaining star data was needed. While not ideal, one method available with the current SDO FSW was to use single shot attitude (SSA) mode to downlink all of the stars (up to 48) in the ST field of view. The major drawback of that approach was SSA mode does not indicate which 9 out of the potentially 48 stars would be used for tracking while in AAD mode. Still, the data could relatively easily be collected and educated guesses could be taken as to which stars were being used for tracking.

The defective pixel table data and the CCD pixel energies are not generated nominally by the star tracker. To obtain that data, commands must be sent to the ST to generate the data and make it available for downlink from one of the diagnostic data subaddresses. Those particular subaddresses are not included in the SDO polling schedule. To downlink any of the diagnostic data, a raw 1553 read command must be sent to that particular subaddress. It is rather straight forward to request generation of the defective pixel table and then use 1553 raw read commands to downlink the data - a total of 17 ground commands is necessary to do so, and an operational procedure was created to implement that sequence of commands.

The data that would be most difficult to obtain was the CCD pixel energies. To obtain that information, or perform a “CCD dump”, another debugging mode, CDI, would need to be used. CDI mode measures the energy level of every pixel within a commanded “window” of the CCD. The CDI command accepts as arguments the starting row, starting column, ending row, and ending column. These “windows” can contain at most 3056 pixels for a single CDI mode command. The ST CCD is composed of 512x512 pixels, but the first and last rows and first and last columns cannot be measured. Figure 17 shows a sketch of the ST CCD and the CDI mode window.

The “CCD dump” cannot be performed while the ST is tracking, so the ST must first be taken out of AAD mode. When the CDI command is issued, the ST automatically transfers from STB to CDI mode to measure the designated pixels and then automatically exits back to STB. Once in STB mode again, the pixel data can be downloaded from the ST, 26 pixels at a time, by reading from one of the diagnostic data subaddresses. The ST will automatically update that subaddress with the next packet of data after each time the subaddress is read. Because the diagnostic data subaddress is not automatically polled by the 1553 bus scheduler, a series of raw read commands are needed to get all of the CCD image data down to the ground. Given the number of pixels that can be imaged in a single CDI mode command and the number of pixels that can be downloaded in a single telemetry packet, the ACS team broke up the “CCD dump” into a total of 85 CDI mode commands that imaged 3046 pixels (6 rows by 507 columns) each, followed by an eighty-sixth CDI mode command (510 rows by 3 columns) that imaged 1530 pixels. Each of the first 85 CDI mode commands had to be followed by 117 raw read commands, containing 26 pixels each. The eighty-sixth CDI mode command was then followed by 85 raw read commands containing 26 pixels each and one raw read command containing the final 22 pixels. To perform a dump of the entire ST CCD based on the above plan would take 10,005 commands to the spacecraft.

Initial testing of the CCD dump procedure indicated that it could take upwards of 16 hours to dump the entire ST CCD if all of the commands were issued from the ground. Having the ST out of use by the onboard attitude estimation FSW for that long would be unacceptable, so the ACS team worked with the FSW team to reduce the time needed to dump the ST CCD. By uploading a series of relative time sequences (RTSs) to issue the CDI mode commands and perform the 1553 raw read commands from the on board FSW, the FSW team was able to reduce the CCD dump time to approximately two hours. Figure 18 shows an image of the ST CCD created using the pixel energy data obtained




Figure 17. Sketch of the ST CCD with a representative CDI window.

from the first CCD dump.

The bright spots in the FOV image denote pixels with higher energy levels. Knowing the ST's attitude when the CCD image was dumped, along with the measured and reference stars reported by the ST, it is possible to determine the possibility of a given pixel being unusually "warm", *i.e.* having a higher background energy level. The data obtained for ST2, along with a similar data set for ST1 was turned over to the ST manufacturer for further analysis. In addition, the SDO AGS team also developed software to process the "CCD dump" data and flag potential warm pixels. More information on the functionality of the analysis software can be found in Ref. 6.

Initial analysis of the ST2 data by both the manufacturer and the AGS team indicated that there were no potential warm pixels near any of the reported stars in the FOV when the low quality index event occurred. A few pixels were found to have higher energy levels, but they were sufficiently far from reported stars as to not have affected the ST




Figure 18. Image of the ST2 CCD using pixel energy data.

solution. However, when the spacecraft was slewed back to that location in space, the same low quality index signature was seen again. As it turns out, that particular attitude is near the bright star, Vega. The current working hypothesis is that Vega could be causing a secondary image on the ST CCD. More discussion of this analysis and hypothesis can be found in Ref. 6.

Because the initial analysis by the vendor and the AGS team did reveal some potentially warm pixels, even if they did not cause the first low quality index event, the decision was made by the FOT, ACS, AGS, and Systems teams to continue to periodically dump the ST CCDs and use the AGS software to analyze the data and trend any potential warm pixels. In the event that the number of potential warm pixels grows over time to the point that the ST solution performance is affected, additional mitigation strategies could then be implemented. More information on additional mitigation strategies can be found in Ref. 6.

## V. Conclusion

Initial checkout of the fine attitude determination sensors on the SDO spacecraft revealed two interesting events. Dependency of the IRU biases on temperature were known prior to launch, but the oscillation of the bias was both unexpected and unacceptable to performance. Though the cause of the oscillation was not able to be determined, the ACS team was able to successfully mitigate the effects of the bias oscillation through the retuning of the Kalman Filter. While the new KF parameters resulted in a slightly noisier solution, the resultant noise is well within performance requirements.

Due to the Herschel ST warm pixel anomaly, the SDO team was not completely unprepared for the inverted “U” shaped quality index. FSW already had a check built in to prevent poor quality ST solutions from corrupting the Kalman Filter. The challenge to the ACS, FSW, and Operations teams was developing the means to bring down the diagnostic data necessary to monitor the development of potential warm pixels on the ST CCD. Through collaboration, the teams were successful in overcoming that challenge, enabling the Operations team to continue to monitor the health of the ST CCDs.

## Acknowledgments

The authors would like to acknowledge Manuel Maldonado, SDO Flight Software Systems engineer for his assistance in writing the ST CCD dump software procedures. The authors would also like to acknowledge Joe Hashmall and Denis Felikson of the Ground Attitude Determination group for their insights into the IRU bias oscillations and their assistance with developing analysis software for monitoring potential warm pixels on the ST.

## References

- <sup>1</sup>Morgenstern, W. M., Starin, S. R., Bourkland, K. L., Liu, K.-C., Mason, P. A. C., and Vess, M. F., “Attitude Control System Design for the Solar Dynamics Observatory,” *AIAA Guidance Navigation and Control Conference*, Portland, OR, August 2011, to be published.
- <sup>2</sup>O’Donnell Jr., J. R., Bourkland, K. L., Hsu, O. C., Liu, K.-C. A., Mason, P. A. C., Morgenstern, W. M., Russo, A. M., Starin, S. R., and Vess, M. F., “Solar Dynamics Observatory Launch and Commissioning,” *34th Annual AAS Guidance and Control Conference*, Univelt, Incorporated, San Diego, California, February 2011.
- <sup>3</sup>Mason, P. and Starin, S. R., “The Effects of Propellant Slosh Dynamics on the Solar Dynamics Observatory,” *AIAA Guidance Navigation and Control Conference*, Portland, OR, August 2011, to be published.
- <sup>4</sup>Vess, M. F., Starin, S. R., and Morgenstern, W. M., “Use of the SDO Pointing Controllers for Instrument Calibration Maneuvers,” *Flight Mechanics Symposium*, Greenbelt, MD, 2005.
- <sup>5</sup>Oort, M., Palomba, M., Procopio, D., Bacchetta, A., Roche, Y., Dugate, D., Pigg, M., Hardacre, S., Seemann, C., and Ochoa, M., “Life in L2: Herschel Early In-Orbit Experience,” *33rd Annual AAS Guidance and Control Conference*, Univelt Incorporated, San Diego, California, February 2010.
- <sup>6</sup>Felikson, D., Ekinci, M., Hashmall, J. A., and Vess, M., “On-Orbit Solar Dynamics Observatory (SDO) Star Tracker Warm Pixel Analysis,” *AIAA Guidance Navigation and Control Conference*, Portland, OR, August 2011, to be published.