

DETERMINING A METHOD OF ENABLING AND DISABLING THE INTEGRAL TORQUE IN THE SDO SCIENCE AND INERTIAL MODE CONTROLLERS

Melissa F. Vess and Scott R. Starin
NASA Goddard Space Flight Center, Greenbelt, MD 20771

ABSTRACT

During design of the SDO Science and Inertial mode PID controllers, the decision was made to disable the integral torque whenever system stability was in question. Three different schemes were developed to determine when to disable or enable the integral torque, and a trade study was performed to determine which scheme to implement. The trade study compared complexity of the control logic, risk of not reenabling the integral gain in time to reject steady-state error, and the amount of integral torque space used. The first scheme calculated a simplified Routh criterion to determine when to disable the integral torque. The second scheme calculates the PD part of the torque and looked to see if that torque would cause actuator saturation. If so, only the PD torque is used. If not, the integral torque is added. Finally, the third scheme compares the attitude and rate errors to limits and disables the integral torque if either of the errors is greater than the limit. Based on the trade study results, the third scheme was selected.

Once it was decided when to disable the integral torque, analysis was performed to determine how to disable the integral torque and whether or not to reset the integrator once the integral torque was reenabled. Three ways to disable the integral torque were investigated: zero the input into the integrator, which causes the integral part of the PID control torque to be held constant; zero the integral torque directly but allow the integrator to continue integrating; or zero the integral torque directly and reset the integrator on integral torque reactivation. The analysis looked at complexity of the control logic, slew time plus settling time between each calibration maneuver step, and ability to reject steady-state error. Based on the results of the analysis, the decision was made to zero the input into the integrator without resetting it. Throughout the analysis, a high fidelity simulation was used to test the various implementation methods.

INTRODUCTION

The Solar Dynamics Observatory (SDO), scheduled to launch in early 2009, will carry a suite of three instruments — the Atmospheric Imaging Assembly (AIA), the Helioseismic & Magnetic Imager (HMI), and the Extreme Ultraviolet (EUV) Variability Experiment (EVE) — designed to provide observations leading to a more complete understanding of the solar dynamics that drive variability in the Earth's environment. The spacecraft's geosynchronous orbit allows for uninterrupted, high-rate downlink of the science data. The spacecraft maintains a fixed attitude relative to the Sun, allowing the instruments to collect a steady stream of solar images.

AIA, built by Lockheed Martin Solar and Astrophysics Laboratory (LMSAL), seeks to investigate the evolution of the Sun's magnetic field through the use of coronal images. AIA also includes four guide telescopes (GT) that are used to drive the instrument's Image Stabilization System (ISS), which prevents high frequency jitter (greater than 10 Hz) from blurring the image data. The GTs are high precision Sun sensors with a fine pointing range of approximately ± 95 arcsec. The data from one of the guide telescopes, called the controlling guide telescope (CGT), is used by the on-board attitude control system (ACS) during Science mode to measure attitude errors relative to the Sun center. HMI, built jointly by LMSAL and Stanford University, is designed to use observations of polarized light to measure the magnetic field and velocity of the solar photosphere. EVE, built by the Laboratory for Atmospherics and Space Physics (LASP) at the University of Colorado, seeks to understand the highly variable solar EUV electromagnetic radiation

and its impacts on the geospace environment. More information about the SDO science instruments can be found through the SDO website.¹

SDO's science mission, in particular HMI's helioseismology, requires twenty-two individual 72-day periods over five years to meet the full science data capture objectives. SDO meets this requirement by providing near-continuous Sun-pointing observations with a few interruptions for necessary maneuvers such as orbit maintenance, momentum management and science instrument calibration. However, the length of these interruptions must be minimized in order to successfully meet the data capture budget.

To ensure the continuing validity of the science data collected by the three instruments, once every three months or so the mission team will command a series of instrument calibration maneuvers sweeping the instrument boresights across the Sun in small steps. For a description of these maneuvers, see Reference 2. Some of these maneuvers are very small, and thus can be performed in Science mode, which uses the GT as the attitude error source.³ The GT has a small fine-pointing field-of-view of approximately ± 95 arcsec. Some of the calibration maneuvers pull the Sun out of the fine-pointing range of the GT, and therefore attitude control during these maneuvers must be performed in Inertial mode, which does not use the GT data.³ Both the Science and Inertial mode controllers, therefore, must provide the required fine pointing and accurate knowledge of the vehicle's attitude relative to the Sun center. These considerations resulted in the decision to adapt a linear proportional-integral-derivative (PID) controller by adjusting gains and adding an attitude limiter to use during Science and Inertial modes [†]. These adjustments ensure that the spacecraft can slew quickly and still achieve steady pointing once on target.

BACKGROUND: SUMMARY OF CONTROLLER STABILITY ANALYSIS

During the analysis of the Science and Inertial mode controller, questions arose about the stability of the controller during slewing maneuvers due to the combination of the integral torque, attitude limit, and actuator saturation. A complete discussion of the subsequent stability analysis performed can be found in Reference 2. The following is not intended as a proof of stability, but is offered as a summary of previous work.

During the analysis, the Routh stability criterion was initially used to assess the possibility of the system going unstable. It showed that because of the attitude and torque saturation, there is a theoretical possibility that the system could go unstable. Further analysis was done to break down the non-linear system into a piece-wise linear system. First, root locus analysis was used to show that with the selected controller gains and in the absence of actuator saturation and an integral torque, attitude limiting alone could not cause an instability in the system. To ensure stability of the described simplified system, the decision was then made to disable the integral torque whenever the attitude error was greater than the attitude limit. However, due to the hardware and software limits on the reaction wheel torques, actuator saturation actually occurs at attitude errors more than an order of magnitude smaller than the attitude limit. Therefore, disabling the integral torque at attitude saturation cannot guarantee total system stability.

On the other hand, torque saturation coupled with an integral torque does not inherently cause system instability; it simply implies that instability is theoretically possible. Further analysis, again using root locus techniques, showed that without attitude saturation, as long as the system output torque is at least 17% of the desired control torque, even in the presence of an integral torque, the system would remain stable. However, at 17% of the desired control torque, there is no stability margin on the system. Project requirements state that there must be at least 6 dB of gain margin on the system.

In summary, based on this analysis, the authors realized there needed to be a scheme for determining exactly when and how to disable the integral torque in order to not just ensure system stability, but also ensure sufficient gain margin on the system. The scheme had to ensure that the integral torque is disabled not only when the attitude error is saturated but also when the attitude error is not saturated but the applied control torque is saturated too much.

[†]The authors wish to acknowledge Dr. Kuo-Chia (Alice) Liu of Goddard Space Flight Center for her help and advice with this analysis effort.

TRADE STUDY ON DISABLING THE INTEGRAL TORQUE

Determining exactly when to disable the integral torque presented an interesting design challenge. The scheme used had to be conservative enough to ensure stability with margin, but not so conservative that the integral torque was not enabled in time to reject steady-state error. Three potential solutions were proposed and a trade study performed to select the best solution based on the following three factors: complexity of the control logic, risk of not reenabling the integral gain in time to reject steady-state error, and the amount of integral torque space utilized. When considering the complexity of the control logic, it is important to note that the SDO controllers are created in wire diagram form using Mathwork's Simulink software and converted to flight software C-code using Mathwork's Real-Time Workshop automated code generation software. Occasionally, implementing a certain design in wire diagram form is more complicated than coding it directly.

Torque Space Description

Although perhaps not directly apparent, the third consideration of the trade study is directly related to the stability margin of the system. Recall that the root locus analysis showed that with an integral torque and in the absence of attitude saturation, the system output torque only had to be 17% of the desired control torque in order for the system to remain stable, but that utilizing that entire torque space would result in no stability margin on the system. The only time the output torque will be less than the desired control torque is if the desired control torque is greater than the actuator saturation limit, τ_{sat} , which for SDO is 0.25 Nm. So for stability with an integral torque, the desired control torque, τ_d , must be less than 1.47 Nm. Figure 1 shows a representation of this stability torque space plotted in the phase plane.

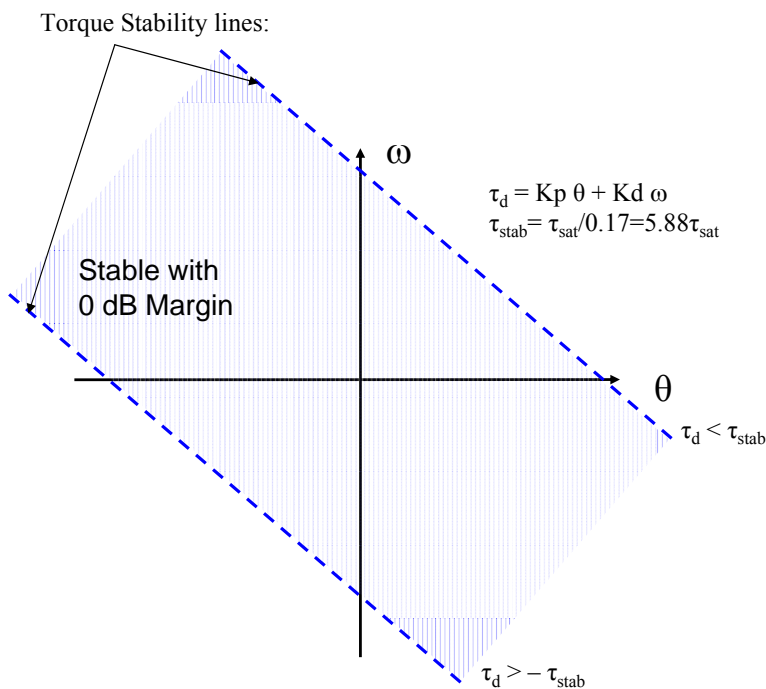


Figure 1: Phase plane plot of maximum integral torque space (no stability margin). Desired torque, τ_d , inside shaded area will result in system stability with the inclusion of the integral torque.

The torque limit lines represent the edges of the allowable desired PD control torque that will ensure system stability. If the desired proportional-derivative (PD) control torque falls inside the shaded area, then the system is stable with the inclusion of the integral torque. If, however, the desired PD control torque falls

outside the shaded area, then the system could potentially be unstable if an integral torque were included. Again, this plot represents the integral torque space assuming no stability margin on the system. In practice, however, the requirement is to have 6 dB of gain margin on the system. Based on the root locus analysis done previously, 6 dB of gain margin means that the output control torque has to be at least 34% of the desired control torque. So, for SDO, τ_d must be less than 0.74 Nm. Figure 2 shows a phase plan representation of this margin torque space.

In this case, the torque limit lines represent the edges of the allowable desired PD control torque that will ensure system stability with adequate margin. So, any scheme used to enable and disable the integral torque can use, at most, the torque space described in Figure 2. Furthermore, the more integral torque space used, the more we take advantage of the available torque capability. However, as will be discussed in the trade study results, that usage has to be traded against the complexity of the control logic.




Figure 2: Phase plane plot of integral torque space with 6 dB of stability margin. Desired torque, τ_d , inside shaded area will result in system stability with 6 dB margin with the inclusion of the integral torque.

Trade Study Options and Results

As was mentioned previously, there were three proposed methods of enabling and disabling the integral torque. The first method proposed using the Routh criterion directly to determine when to disable the integral torque. The second method proposed looking at the desired PD torque and attitude error to determine when to disable the integral torque. And finally, the third method compared the attitude and rate errors to limits to determine when to disable the integral torque, ensuring that the output control torque over the desired control torque falls within the margin torque space.

Option One: Using the Routh Criterion

In Reference 2, for the Science and Inertial mode controllers, the simplified Routh criterion is:

$$k_i < k_p * k_d * k_t \quad (1)$$

where the structural filter, one-cycle delay, gyro and wheel dynamics, and the plant flexible modes have been neglected and k_i , k_p , k_d , and k_t are the integral, proportional, derivative, and torque gains, respectively. For Science mode X-axis and Inertial mode $k_i = 0.005129 \text{ Nm/kg} \cdot \text{m}^2 \cdot \text{s} \cdot \text{rad}$, $k_p = 0.07192 \text{ Nm/kg} \cdot \text{m}^2 \cdot \text{s}$, and $k_d = 0.467055 \text{ Nms/kg} \cdot \text{m}^2 \cdot \text{rad}$. The torque gain, k_t is the ratio of the output torque over the desired torque and represents the amount of torque scaling present.

This first method of enabling and disabling the integral torque proposes calculating $k_p * k_d * k_t$ each cycle and comparing it to k_i . If k_i is greater than $k_p * k_d * k_t$, the integral torque will be disabled. Otherwise, the integral torque will be enabled. The main advantage of this method is that it theoretically takes advantage of all of the integral torque stability space. In fact, implementing the enabling/disabling using Equation 1 directly does not ensure sufficient gain margin, and further analysis would be needed to ensure the gain margin meets the 6 dB requirement.

There are three main disadvantages to using the Routh criterion. The first is that, as was mentioned above, further analysis would have to be done to modify Equation 1 to ensure the 6 dB of gain margin. The second is that this method relies on using the equation for the simplified Routh criterion, which neglects the structural filter, one-cycle delay, gyro and wheel dynamics, and the plant flexible modes. Because these dynamics exist in the real spacecraft, passing a test based on Equation 1 cannot guarantee stability of the complete system.

The final drawback to this method is that it would be complicated to implement. In order to calculate k_t , both the PD and PID torque commands would have to be calculated to get the ratio of the two. In addition, because of the attitude limiter, k_p is not effectively constant. Similarly to k_t , the effective k_p would have to be calculated as the ratio of the limited attitude over the input attitude times the attitude gain. This multitude of calculations would make coding and testing the software complicated, expensive, and susceptible to error.

Option Two: Combining PD Torque and Attitude Saturation

The second proposed method of disabling the integral torque involves calculating the PD torque and checking to see if the attitude is saturated. If the PD torque is less than the wheel saturation limit and if the attitude error is not saturated, then calculate the integral torque and add it to the PD torque. Otherwise, if the PD torque is greater than the wheel saturation limit or if the attitude error is saturated, then use just the PD torque. This method ensures that neither the attitude saturation nor the actuator saturation can induce instability in the system. The main advantage of this method is that it utilizes a lot of the available integral torque margin space, but still ensures sufficient stability margin. Figure 3 shows a sketch of the integral torque space utilized by Option Two.




Figure 3: Phase plane plot of integral torque space utilized by Option Two.

In the plot, θ_a represents the attitude saturation limit. Notice that the plot only shows the torque saturation lines, τ_{sat} . Recall from Figure 2 that the outer bounds of the margin space fall at $2.94\tau_{sat}$, so Option Two's integral torque space utilized falls within the margin space. Theoretically, Option Two could be adjusted to utilize even more of the margin space by comparing the PD torque to $2.94\tau_{sat}$, in which case the switching lines of Option Two would be identical to those in Figure 2.

The only major drawback to this method is that it would be complicated to implement. While easier than Option One, Option Two still requires several calculations and complicates the controller wire diagram. First, just the PD torque has to be calculated. Then the attitude error has to be compared against the attitude saturation limit. Then there has to be logic to determine if those two elements pass or fail. Then, if both of those elements pass, the integral torque has to be calculated and added to the PD torque. Again, this number of calculations would make coding and testing the software complicated and expensive, and potentially error prone.

Option Three: Using Attitude and Rate Limits

The third and final proposed method of integral torque disabling involves comparing the attitude and rate errors against limits and disabling the integral torque if either the attitude error or the rate error falls outside the limits. The limits are determined such that the PD torque calculated using the attitude and rate limits falls inside the torque margin space. Figure 4 shows a sketch of the integral torque space utilized by Option Three.

As with Figure 3, this sketch shows only the torque saturation lines, which, as was stated above, fall inside the torque margin space. In the plot, θ_{lim} and ω_{lim} represent the attitude and rate limits, respectively, used to disable the integral torque. Notice that these limits form a box in the phase plane, which will be




Figure 4: Phase plane plot of integral torque space utilized by Option Three.

called the “integral torque box”. As in Figure 3, θ_a represents the attitude saturation limit. The θ_a lines are shown to emphasize that θ_{lim} will be smaller than the attitude saturation limit.

The major advantage to this method is its ease of implementation. The only additional calculations that have to be done are comparing the attitude and rate errors to θ_{lim} and ω_{lim} , respectively. If both the attitude and rate errors fall within the integral torque box, then the integral torque is calculated, otherwise the PD torque is calculated. The major disadvantage to this method is that it is very conservative in its use of the integral torque margin space. Of the three proposed methods, it uses the least amount of that space. In addition, finding appropriate values for θ_{lim} and ω_{lim} can be somewhat tricky. The limits have to be wide enough to ensure that the system can get back within the integral torque box so that steady-state error can be rejected. At the same time, the limits cannot be too wide or the 6 dB gain margin requirement will not be met.

Despite its conservatism and somewhat complex limit selection, Option Three was selected as the method of enabling and disabling the integral torque. Initially, θ_{lim} and ω_{lim} were selected such that:

$$\theta_{lim}K_{p,max} + \omega_{lim}K_{d,max} < \tau_{sat} \quad (2)$$

where $K_{p,max}$ and $K_{d,max}$ are the maximum attitude and rate gains, respectively. The SDO controller gains were designed independent of the spacecraft moments of inertia, therefore the total gains are proportional to the moments of inertia. At the time of analysis, $K_p = [126, 213, 166]$ Nm/rad and $K_d = [780, 1317, 1023]$ Nm/rad/sec. For SDO, the axis of maximum moment of inertia is the Y-axis, therefore $K_{p,max}$ and $K_{d,max}$ were equal to 213 Nm/rad and 1317 Nm/rad/sec, respectively. The saturation torque, τ_{sat} at the time of the analysis was 0.2 Nm. The main consideration in setting ω_{lim} was that the least significant bit of the IRU results in a rate resolution of about 1.2 rad/sec. At steady-state, when the spacecraft has no rate, noise in the system will cause the IRU measured spacecraft rate to fluctuate between

$\pm 1.2e - 5$ rad/sec. ω_{lim} should therefore be set to at least $1.2e-5$ rad/sec. Initially, a value of $2.62e-5$ rad/sec was chosen for ω_{lim} , which is slightly more than twice the IRU resolution. Once ω_{lim} was set, a value for θ_{lim} could be backed out using Equation 2. The θ_{lim} chosen was $6.98e-4$ rad, which is equivalent to 144 asec. The selected values for θ_{lim} and ω_{lim} resulted in a maximum torque of 0.18 Nm, which is less than τ_{sat} and therefore stability is ensured.

INTEGRAL TORQUE RESET ANALYSIS

Once the method of determining when to disable and enable the integral torque was selected, the next step was to determine how to disable and enable the integral torque. Like determining when to disable, determining how to disable the integral torque also posed an interesting design problem. The problem could be broken down into two parts; how exactly to go about disabling the integral torque and whether or not the integrator should be reset upon exiting or entering the integral torque box.

Early in the redesign process, the decision was made that the easiest way, in terms of implementation, to disable the integral torque would be to zero the input into the integrator whenever the attitude or rate errors fell outside the integral torque box limits. Zeroing the input into the integrator holds the integral torque constant while outside the integral torque box, thereby acting as a constant disturbance torque on the system. Other options considered, but rejected, were zeroing the integral torque directly but allowing the integrator to continue integrating while outside of the integral torque box and zeroing the integral torque directly and resetting the integrator upon reentry to the integral torque box. The first reason these methods were rejected was because zeroing the integral torque directly is more complicated to implement in the controller wire diagram than zeroing the input into the integrator. In addition, during spacecraft slews the first method can result in saturation of the integrator due to large attitude errors. The purpose of the integral torque term is to compensate for external disturbance and wheel drag torques during steady-state pointing, which in turn improves controller performance by improving steady-state pointing accuracy. Allowing the integrator to continue integrating and potentially saturate during slews adds no benefit to steady-state controller performance.

When the selected method of disabling the integral torque was presented at a peer review, the concern was voiced that by not resetting the integrator upon leaving the integral torque box, the held integrator torque, which acts as a disturbance torque, could be large enough that the controller would not be able to get back inside the integral torque box and a steady-state attitude error would result. The authors then decided to do further analysis on the proposed method, to determine the possibility of not being able to reenter the integral torque box and also to look into the possibility of resetting the integrator any time the integral torque box is exited. In addition to the ability to reject steady-state error by reentering the integral torque box, the analysis also had to consider complexity of the control logic and slew time plus settling time between steps of the instrument calibration maneuvers.

The control torque, τ_c is calculated as follows:

$$\tau_{c,ss} = K_p\theta + K_d\dot{\theta} + K_i \int \theta dt \quad (3)$$

where θ is the attitude error and K_p , K_d , and K_i are the proportional, derivative, and integral gains, respectively, and include multiplication by the spacecraft inertia, *e. g.* $K_p = k_p * \mathbf{I}$. When the system is at steady-state, the attitude and rate errors are zero, and, as stated above, the integral term acts to compensate for the external disturbance and wheel drag torques, τ_{ext} and τ_{drag} , respectively.

$$\tau_{c,ss} = -(\tau_{ext} + \tau_{drag}) = K_i \left(\int \theta dt \right)_{ss} \quad (4)$$

When the attitude and/or rate errors exceed the integral torque box limits, for example when an attitude slew is commanded, the input into the integrator is zeroed and the value of the integral torque at that instant

is maintained. Let that torque be called τ_i . Then, during the slew, the control torque can be calculated as follows:

$$\tau_c = K_p\theta + K_d\dot{\theta} + \tau_i \quad (5)$$

At the end of the slew, the rate error will go to zero, and the above equation can be rewritten as follows:

$$\tau_c = K_p\theta_{ss} + \tau_i \quad (6)$$

where θ_{ss} is the steady-state attitude error. As before, when the system is at steady-state, the controller acts to compensate for the external disturbance and wheel drag torques:

$$-(\tau_{ext} + \tau_{drag}) = K_p\theta_{ss} + \tau_i \quad (7)$$

In order for θ_{ss} to go to zero, we must get back inside the integral torque box. In order for that to happen, the following inequality must be true:

$$|\tau_{ext} + \tau_{drag} + \tau_i| < |K_p\theta_{lim}| \quad (8)$$

where θ_{lim} is the attitude limit of the integral torque box, which was originally set to 6.98e-4 rad. Consider the case when the integrator is saturated just prior to leaving the integral torque box. The integrator saturation limit is 0.01 rad-sec. Because of changes in the spacecraft inertias, the gains used in this analysis are different from those used in earlier analysis. The integral gain, including spacecraft inertia, used in the analysis is [11.3, 20.3, 16.8] Nm/rad-sec. The proportional gain, including spacecraft inertia used in the analysis, is [159, 285, 235] Nm/rad. The spacecraft inertia used for the analysis assumes that the spacecraft has full fuel tanks, *i. e.* the inertias are at their greatest values. Using the Y-axis gains as the worst case and plugging into the above equation yields:

$$\tau_{ext} + \tau_{drag} + 0.203 \not< 0.199 \quad (9)$$

which implies that if the integrator is at its saturation limit, the controller may, in the worst case, be unable to get the spacecraft back inside the integrator torque box. The question then is what is the likelihood of reaching the integrator saturation limit. As shown in Equation 2, the integrator torque when the spacecraft is at steady-state is equal in magnitude to the sum of the external disturbance and wheel drag torques. For the SDO orbit, the wheel drag torque actually dominates the external disturbance torques. The high fidelity simulation (HiFi), which includes a wheel drag model, shows that it takes approximately 130 Nms of momentum in the Y-axis before the integration saturation limit is met. This amount of momentum translates into almost 78 Nms of momentum in each of two wheels. The estimated worst-case momentum build-up between momentum unloadings is a total of approximately 15 Nms, which puts about 8.66 Nms in each of two wheels. So, failing to get back inside the integral torque box due to integrator saturation cannot happen during expected nominal operations.

Another case to consider when determining the possibility of not reentering the integral torque box is when the wheel momentum is distributed such that all the system momentum is oriented in the Y-axis (axis of largest moment of inertia) and the spacecraft does a 180-degree roll about the X-axis. Such a maneuver is possible during a contingency high gain antenna handover maneuver. At the beginning of the maneuver, τ_i is approximately equal in magnitude to, and acts to counteract, the external and wheel drag torques. Let the sum of those terms at the beginning of the slew be τ_b . At the end of the maneuver, the system momentum vector has rotated 180 degrees and is now pointed along the -Y-axis. The wheel momentum has likewise

switched directions, and thus the sum of the external and wheel drag torques at the end of the maneuver, τ_e , is equal in magnitude and opposite in direction to τ_b . Going back to Equation 8,

$$\begin{aligned} |\tau_{ext} + \tau_{drag} + \tau_i| &< |K_p \theta_{lim}| \\ |\tau_e + \tau_i| = |\tau_e - \tau_b| &< |K_p \theta_{lim}| \\ |-2\tau_b| &< |K_p \theta_{lim}| \end{aligned} \quad (10)$$

Depending on the amount of momentum in the wheels, it is possible that the inequality in Equation 10 will not be met. Using the HiFi, if the system momentum in the body frame is equal to $[0, 64, 0]$ Nms ($[-36.95, 0, 36.95, 0]$ Nms in the wheels³), and the spacecraft performs a 180 degree roll about the X-axis, then the controller can barely get the spacecraft back inside the integral torque box. Any more momentum in the Y-axis, and the controller cannot get the spacecraft back inside the integral torque box. If the integrator is reset, then τ_i is zero, and the controller has no problems getting back inside the integral torque box. While that amount of momentum in the wheels is still outside the predicted nominal operations, it is possible that for long slews the integrator will need to be reset. While the likelihood is that resetting the integrator would not be necessary for the short instrument calibration slews, from an operational simplicity stand point, the question becomes whether the integrator be reset *anytime* the integral torque box is exited.

To answer this question, we considered the instrument calibration maneuvers, which consist of a series of small slews (on the order of 5 arcsec – 9000 arcsec) with fast slew time requirements (2 – 5 minutes). Using the HiFi, the spacecraft was started with approximately $[1, 34, 5]$ Nms of momentum in the spacecraft body frame and performed the EVE Cruciform calibration maneuver.² If the integrator was not reset each time the integral torque box was exited, the spacecraft had no problems performing the slews in their allotted time and reentering the integral torque box. If, however, the integrator was reset each time the integral torque box was exited, the spacecraft could not perform the slews within the allotted time.

Recall from Equation 4 that the integral torque acts to compensate for the external and drag torques. Because the momentum in the wheels is relatively constant over short time frames, the integral torque is also roughly constant. The value of the integrator, therefore, is proportional to the amount of momentum in the wheels– the higher the momentum, the larger the integrator values. Every time the integrator is reset, it takes time for the integrator to re-converge to its steady-state value. The time constant for the integrator is equal to K_p/K_i , which is approximately 14 sec. The controller is simultaneously trying to re-converge the integrator and zero the attitude and rate errors. The result is that it takes longer to perform each slew. Given the time constant of the integrator and the short slew times, there is not enough time between each slew to reset the integrator.

Based on the above analysis, the decision was made not to reset the integrator automatically upon exiting the integral torque box, but to have a flight software command available to reset the integrator should the need arise.

SELECTION OF INTEGRATOR TORQUE BOX LIMITS

Once the decision was made not to reset the integrator automatically upon exiting the integral torque box during nominal mission operations, the authors decided to revisit the selected torque box limits, θ_{lim} and ω_{lim} to see if it was possible to increase those limits to further reduce the possibility of needing to reset the integrator manually. Using Equation 7, if the spacecraft does not get back inside the integral torque box, the worst expected steady-state pointing error is:

$$\theta_{ss} = \frac{\tau_{ext} + \tau_{drag} + \tau_i}{K_p} \quad (11)$$

However, as was seen in Equation 10, the worst case torque magnitude is equal to two times τ_b , which is the sum of τ_{ext} and τ_{drag} , so:

$$\theta_{ss} = \frac{2\tau_b}{K_p} \quad (12)$$

To be able to get back inside the integral torque box, θ_{lim} must be greater than θ_{ss} . At this point, assume a θ_{lim} larger than θ_{ss} and assume a rate limit, ω_{lim} small enough to indicate slewing has stopped. Calculate the maximum PD torque seen when both θ_{lim} and ω_{lim} are reached:

$$\tau_{lim} = K_p\theta_{lim} + K_d\omega_{lim} \quad (13)$$

where τ_{lim} is the maximum the PD torque command can be while still including the integral term. Recall from Figure 2 that to ensure sufficient margin, τ_{lim} has to be less than $\tau_{sat}/0.34$, or 0.74 Nm. Plugging into the above equation and using the original integral torque box limits, τ_{lim} is calculated to be [0.138, 0.247, 0.204] Nm. As was done before, the spacecraft inertia was assumed to include full fuel tanks. So, as can be seen from the above calculation, it is possible to increase θ_{lim} and ω_{lim} . The authors decided on a value of 250 arcsec or 1.212e-3 rad for θ_{lim} and 4.0e-5 rad/sec for ω_{lim} . These limit values result in a τ_{lim} equal to [0.234, 0.419, 0.346] Nm. Even for the maximum axis (Y), τ_{lim} is still less than the margin stability value of 0.74 Nm. Theoretically, θ_{lim} and ω_{lim} could be increased further, but the authors would like to retain the additional margin. If issues with returning to the integral torque box arise on orbit, then the limits can be increased.

CONCLUSION

Through analysis and simulation, the authors instituted a method of enabling and disabling the integral torque portion of the Science and Inertial mode PID controllers to ensure system stability. The method compares the absolute value of each component of the attitude and rate errors to a predetermined limit. If the attitude and rate errors are below the limit, the integral torque is added to the proportional and derivative torques, otherwise the input into the integrator is zeroed such that the integral torque remains constant and acts as a disturbance torque on the system. The attitude and rate limits were chosen such that at the limits of the integral torque box, the controller is still guaranteed stable with sufficient gain margin. The authors decided against resetting the integrator upon exit from the integral torque box because analysis showed that it would take too long for the integrator to reconverge between instrument calibration slews. Given the current reaction wheel torque limits and expected momentum build-up, even with the worst case expected integral value acting as a disturbance torque, the system still has sufficient torque authority to return to the integral torque box and reject steady-state errors. Should issues arise on orbit, a flight software command does exist to reset the integrator, or the integral torque box limits can be extended.

REFERENCES

- [1] Solar Dynamics Observatory Home Page. <http://sdo.gsfc.nasa.gov>. responsible official: Liz Citrin, Goddard Space Flight Center, August 2007.
- [2] Melissa F. Vess, Scott R. Starin, and Wendy M. Morgenstern. Use of the SDO Pointing Controllers for Instrument Calibration Maneuvers. *Flight Mechanics Symposium*, 2005.
- [3] 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.