

ATTITUDE GROUND SYSTEM (AGS) FOR THE MAGNETOSPHERIC MULTISCALE (MMS) MISSION

Juan C. Raymond,⁽¹⁾ Joseph E. Sedlak,⁽²⁾ and Babak Vint⁽²⁾

⁽¹⁾ NASA/Goddard Space Flight Center, Code 591, Greenbelt, MD 20771, USA, juan.c.raymond@nasa.gov

⁽²⁾ a.i. solutions, Inc., 4500 Forbes Blvd., Suite 300, Lanham, MD 20706, USA
joseph.sedlak@ai-solutions.com and babak.vint@ai-solutions.com

Abstract: *The Magnetospheric Multiscale (MMS) mission is a Solar-Terrestrial Probe mission consisting of four identically instrumented spin-stabilized spacecraft flying in an adjustable pyramid-like formation around the Earth. The formation of the MMS spacecraft allows for three-dimensional study of the phenomenon of magnetic reconnection, which is the primary objective of the mission. The MMS spacecraft were launched early on March 13, 2015 GMT. Due to the challenging and very constricted attitude and orbit requirements for performing the science, as well as the need to maintain the spacecraft formation, multiple ground functionalities were designed to support the mission. These functionalities were incorporated into a ground system known as the Attitude Ground System (AGS).*

Various AGS configurations have been used widely to support a variety of three-axis-stabilized and spin-stabilized spacecraft missions within the NASA Goddard Space Flight Center (GSFC). The original MMS operational concept required the AGS to perform highly accurate predictions of the effects of environmental disturbances on the spacecraft orientation and to plan the attitude maneuvers necessary to stay within the science attitude tolerance. The orbit adjustment requirements for formation control drove the need also to perform calibrations that have never been done before in support of NASA GSFC missions. The MMS mission required support analysts to provide fast and accurately calibrated values of the inertia tensor, center of mass, and accelerometer bias for each MMS spacecraft.

During early design of the AGS functionalities, a Kalman filter for estimating the attitude, body rates, center of mass, and accelerometer bias, using only star tracker and accelerometer measurements, was heavily analyzed. A set of six distinct filters was evaluated and considered for estimating the spacecraft attitude and body rates using star tracker data only. Four of the six filters are closely related and were compared during support of the Time History of Events and Macroscale Interactions during Substorms (THEMIS) and Space Technology-5 (ST-5) missions. These analyses exposed high dependency and sensitivity on the knowledge of the spacecraft inertia tensor for both body rates and accelerometer bias estimation. The conclusion of the analysis led to the design of an inertia tensor calibration technique using only star tracker data. The second most important result of the analysis was the design of two separate Kalman filters to estimate the spacecraft attitude and body rates and the accelerometer bias instead of a single combined filter. In this paper, the calibration results of the mass properties, as well as the performance of the spacecraft attitude and body rates filters using flight data are presented and compared against the mission requirements.

Keywords: *MMS, Kalman filters, SpinKF, calibration of mass properties, inertia tensor calibration.*

1. Introduction

This paper describes the Attitude Ground System (AGS) design and implementation used for support of the Magnetospheric MultiScale (MMS) mission. The AGS is one of the components within the mission operations center providing flight dynamics support to the mission. The initial design for the AGS was described in [1]. MMS was launched on March 13, 2015, as the fourth mission of the NASA's Solar Terrestrial Probe (STP) program.

This work was supported in part by the National Aeronautics and Space Administration (NASA)/Goddard Space Flight Center (GSFC), Greenbelt, MD, USA, Contract NNG14VC09C.

The MMS mission consists of four identical spacecraft flying in a tetrahedral formation. The mission is divided into two phases, each with a different type of eccentric Earth orbit. The first phase is designed for an orbit of 1.2×12 Earth radii with a period of roughly one day to study the dayside of the magnetopause. The second phase has an orbit of 1.2×25 Earth radii with a period of approximately 3 days to study the nightside magnetotail. By varying the size of the tetrahedron, the orbital semi-major axis, and the eccentricity, and by making use of the changing solar phase, the geometry allows for the study of both bow shock and magnetotail plasma physics; including acceleration, reconnection, and turbulence.

Each of the four MMS spacecraft spins at approximately three revolutions per minute (rpm), with the spin axis tipped off the ecliptic north pole by approximately 2.5 deg toward the sun-line. This spin rate allows for maintenance of the necessary tension in the four radial 60-meter wire booms carrying some of the science instruments. There also are two axial booms and two radial magnetometer booms. The onboard Attitude Control System (ACS) uses a star tracker system (STS) with four camera heads, a slit-type digital Sun sensor (DSS), an accelerometer, and GPS receivers for attitude and precise orbit determination. Thrusters are used for attitude and orbit maneuver control.

The AGS is responsible for determining and predicting the spacecraft attitude, providing support to the mission operations team to keep the spin axis orientation within the target tolerance. The AGS is also essential for validation of the onboard attitude and body rate estimates, sensor interference predictions, and sensor and inertia tensor calibrations to improve the accuracy of the ground and onboard attitude solutions.

2. AGS Overview

This section describes the AGS functionalities required and designed for MMS support. This paper will not discuss the input and output products to perform the underlying functions of the MMS AGS, which mainly include:

- Estimation and validation of the spacecraft attitude and body rates using an extended Kalman filter.
- Attitude and sensor interference prediction with the use of a smart targeting algorithm.
- Maneuver planning and scheduling support to stay within the science attitude requirements.
- Calibration of mass properties (inertia tensor and center of mass) to improve ground and flight attitude determination solutions.
- Sensor alignment calibrations, such as DSS to Star Sensor and relative Star Sensor Camera Head alignments.

Although the core of the AGS could either be the multi-mission three-axis stabilized spacecraft (MTASS) system [2], or the multi-mission spin stabilized spacecraft (MSASS) system, the former was selected as the baseline for the MMS AGS with features and capabilities added to support spin-stabilized spacecraft and satisfy MMS mission requirements. MMS spacecraft are spinners with an autonomous quaternion-output star tracker system (STS) that make possible the determination of full three-axis attitudes, instead of just the single spin-axis direction.

This wide variety of features place the AGS at the heart of a powerful and invaluable ground support system that has been used for nominal operations, calibrations, data analyses, as well as anomaly resolution of over two dozen missions. The oldest parts of the AGS were written in FORTRAN in the 1970s and 1980s. The generalized multi-mission version [2] was created in the early 1990s and was used to support, e.g., the UARS, EUVE, and RXTE missions. The AGS software was ported to MATLAB in the late 1990s.

3. Ground Attitude Determination and Dependencies

The MMS spacecraft attitude requirement is 0.1 deg (3σ) per axis for both phases of the mission. The resulting history of the spacecraft attitude and body rates is called the definitive attitude. This definitive attitude is generated using quaternion solutions from each of the four sensor heads output by the STS at 0.5 to 4 Hz, depending on the telemetry downlink rate. The STS heads are mounted in pairs on two separate stable optical benches and are oriented roughly 10 deg off the body $-Z$ axis. This orientation was chosen to minimize star motion through the fields of view (FOVs) and to avoid interference from an axial sensor boom deployed along the $-Z$ axis direction. An alignment transformation is applied to the output from each STS camera head so that the final output quaternions represent the geocentric inertial frame (GCI)-to-spacecraft body frame transformation.

3.1. Star Tracker Alignment Calibration Dependency. The pre-launch alignments of the Star Tracker System (STS) heads were measured by optical methods using reflective cubes on the benches. In-flight calibration of the STS heads only corrects for the relative misalignment and not for absolute alignments in the spacecraft body frame. The mean alignment of the four STS heads defines the body frame on orbit. Therefore, all four heads should provide the same relative attitude solutions to the body frame once they are calibrated.

Although the relative alignment of two heads on the same bench remains more stable than two heads on separate benches, calibration of all four head alignments has been performed multiple times during MMS launch and early orbit (L&EO) operations to correct for any shifts due to launch shock, release of gravitational stress, and change of thermal environment (the latter being by far the most significant).

The AGS calibration utility uses the ALIQUEST method [3]. This attitude-dependent method is based on the QUEST quaternion estimation algorithm [4] to solve for the misalignment using a reference attitude. The QUEST algorithm is a well-known, efficient, and reliable method to determine the attitude by minimizing the loss function, L , as a function of the attitude matrix, A ,

$$L(A) = \sum_i \left[\vec{v}_i^{body} - A \vec{v}_i^{ref} \right]^2, \quad (1)$$

where vector \vec{v}_i^{body} is the observation unit vector for sensor i expressed in the body frame, \vec{v}_i^{ref} is the corresponding inertial frame reference vector, and index i runs over all sensors available at a given time. The matrix A is the “single-frame” attitude estimate at that time. Similarly, the attitude-dependent alignment estimation problem can be cast in a parallel form. That is, determine the misalignment for a given sensor by minimizing the loss function

$$L(O) = \sum_j \left[O \vec{v}_j^{body} - A_j \vec{v}_j^{ref} \right]^2 \quad (2)$$

as a function of the orthogonal misalignment matrix, O , where subscript j is a time index running over all valid observations, and A_j is the known attitude history. The vector \vec{v}_j^{body} is the observation unit vector for the given sensor, rotated to the body frame using the nominal alignment, N_o , and any a priori misalignment, M_o . Once the correction matrix, O , is determined, the new misalignment is usually expressed in the sensor frame as

$$M = N_o^{-1} O N_o M_o, \quad (3)$$

where the inverse of N_o is used here rather than the matrix transpose to allow for possible nonorthogonality. It is clear Eqs. 1 and 2 can be minimized using the same method. Choosing the QUEST algorithm [4] to solve Eq. 2 results in the ALIQUEST utility.

For the MMS STS calibration, the reference attitude is generated using approximately the same number of pre-calibration observations from each STS head, all with the same weight. This makes the reference attitude, in effect, an average over all four misalignments. After the calibration, the

alignments of the four STS heads is adjusted to agree with this reference attitude. Thus, when the calibrated STS data is used in an attitude filter, the STS residuals will be reduced but the attitude solution should be unchanged. This means that ALIQUEST calibrations can be performed at any time during the mission without introducing any discontinuity in the definitive attitude histories.

3.2. Definitive Attitude with Star Tracker Data. For all the time spans with valid STS data, the AGS produces a definitive attitude using an estimator referred to as the spinning spacecraft Kalman filter or SpinKF. Six distinct extended Kalman filters were studied for suitability to support the MMS mission. These filters were analyzed with THEMIS and ST-5 data, as well as with high fidelity MMS simulations of the flexible wire booms and other appendages. The chosen filter has been used operationally for THEMIS and ST-5. This filter is described in [5] and [6] and is conceptually related to the “standard” Lefferts, Markley, and Shuster quaternion and rate filter described in [7]. Regardless of the filter chosen, all of them exposed high dependency and sensitivity on the knowledge of the spacecraft inertia tensor for all attitude, and body rates.

The SpinKF’s state vector avoids using the quaternion; instead, it uses a seven-parameter angular momentum-based representation [8]. The seven state vector elements are the angular momentum components in the inertial and body frame, and the spin phase angle. These parameters are subject to the constraint that the magnitude of the angular momentum is the same in the inertial and body frames, just as the standard quaternion and rate filter has seven components with the constraint that the quaternion be normalized.

Most spinning spacecraft, like MMS, do not carry rate-sensing gyros. As a consequence, the SpinKF integrates Euler’s equations of motion to propagate the state between observations. The result of this integration depends highly on the inertia tensor as attitude and transverse body rate solutions lose significant accuracy if knowledge of the inertia tensor is poor. Errors in our knowledge of the inertia tensor yield a non-zero mean in both the attitude residuals and the transverse body rates, over and above those offsets expected from any actual misalignment of the inertia major principal axis and the body Z-axis. These errors, likewise, increase the magnitude or peak-to-peak oscillations in the estimates of the transverse body rates. For this reason, inertia tensor calibration was included in the AGS suite of utilities.

Despite this dependency on the inertia tensor, the numerical integration is more accurate for a set of parameters having less variation, as compared to the quaternion. The angular momentum-based representation [8] has only a single component that changes quickly, the spin phase. This component increases only linearly (mod 2π). The quaternion, on the other hand, varies sinusoidally and with large amplitude for spinners.

3.3. Inertia Tensor Calibration Dependency. The inertia tensor calibration was divided into two parts: major principal axis (MPA) and inertia ratio estimation. Due to the fact that the inertia ratio was only observable very early in the mission, when the MMS spacecraft was in a rigid-body configuration before the booms were deployed, only the MPA correction is discussed in this paper. The inertia ratio calibration is described in [1]. The major principal axis calibration is designed to provide improved knowledge of the direction of the MPA in the body frame.

The algorithm uses the STS measurements to estimate the direction of the angular momentum vector in the inertial frame. By averaging over many spin periods and boom vibration periods, the estimated angular momentum can be co-aligned with the major principal axis of inertia. This provides enough information to determine the direction of the MPA in the body (details are given below and are described in [1]). The a priori inertia tensor is then rotated to agree with this estimated principal axis, without making any other changes to the inertia. This process can be iterated with the STS head alignment calibration because any misalignments falsely introduce a coning correction in the inertia tensor. Coning is understood as the misalignment of the MPA from the body Z-axis while nutation is the offset of the angular momentum vector in the body frame from the MPA.

First, the attitude history is determined from a set of STS measurements. From this, the direction of the body Z-axis, \hat{Z}_I , is computed, where the subscript I here implies the vector is referenced to an inertial reference frame, and the caret indicates unit vector. The angular momentum direction, \hat{L}_I , may be found by calculating the vector around which the Z-axis rotates in inertial space. If there is no nutation, then \hat{Z}_I will rotate around \hat{L}_I on a circular cone. If there is nutation then the motion will be more complicated as \hat{Z}_I will rotate around the MPA, \hat{P}_I , which itself will be rotating around \hat{L}_I . Given enough data and an integral number of nutation periods, \hat{L}_I may be found in either case by taking the time-average of the motion of \hat{Z}_I .

Next, this estimate of the angular momentum direction \hat{L}_I is converted to the body frame using the known attitude history to yield a set of unit vectors, \hat{L}_B , where the subscript B here implies the vector is referenced to the body frame. Then, \hat{L}_B is averaged over the entire span of nutation periods and re-unitized to obtain \bar{L}_B , where the overbar indicates the mean value.

If there is no nutation, all the \hat{L}_B vectors will be parallel (to within the attitude noise) and nearly equal to \bar{L}_B . If there is nutation, \hat{L}_B will sweep out an elliptical cone around \bar{L}_B in the body frame. In either case, the MPA in the body frame, \hat{P}_B , is approximated by the average \bar{L}_B .

Finally, the utility computes the direction cosine matrix (DCM) [9] that rotates the a priori MPA, \hat{P}_B^{prior} , into \hat{P}_B . This DCM is used in a similarity transform on the a priori inertia tensor, J_{prior} , to yield a new tensor whose MPA is the one estimated. That is,

$$\varphi = \cos^{-1}(\hat{P}_B^{prior} \cdot \hat{P}_B), \quad (4)$$

$$\hat{e} = \frac{\hat{P}_B^{prior} \times \hat{P}_B}{|\hat{P}_B^{prior} \times \hat{P}_B|}, \quad (5)$$

$$M = \cos \varphi I_3 + (1 - \cos \varphi) \hat{e} \hat{e}^T - \sin \varphi [\hat{e} \times], \quad (6)$$

$$J_{new} = M^T J_{prior} M, \quad (7)$$

where the caret indicates unit vector and where $[\hat{e} \times]$ is the skew-symmetric cross-product operator

$$[\hat{e} \times] = \begin{bmatrix} 0 & -e_z & e_y \\ e_z & 0 & -e_x \\ -e_y & e_x & 0 \end{bmatrix}. \quad (8)$$

Performing this estimation does not give any information about the directions of the inertia tensor's two transverse principal axes. By using a minimal rotation, M , from the a priori principal axis to \hat{P}_B , the original directions of the transverse principal axes are affected very little.

The reference attitude for this calibration comes from the SpinKF solution (although raw STS attitudes could be used). If the SpinKF is used, there will be some dependence on the a priori inertia tensor. Once a new tensor is obtained from Eq. 7, the filter can be rerun and the process iterated. Convergence has been found to be very rapid; two or three iterations typically are sufficient to determine the principal axis to within less than an arc-second.

Despite the fact that the accuracy of the inertia tensor estimate depends on STS noise and STS head misalignments, test results with simulated data showed very good results even in the presence of coning, nutation, and boom oscillations. Tests since launch with flight data show very consistent inertia estimates with different data sets, with the MPA varying by less than 0.001 deg.

3.4. Definitive Attitude with Star Tracker Data Gaps. If there are any gaps in the STS data, the attitude will be solved for with the assumption that the angular momentum is constant in both inertial and body frames. When filling the gaps, the MPA is aligned with the angular momentum direction (e.g., the mean angular momentum estimated from the data sets before and after the gaps). The spin phase in the gaps is computed using the Sun pulse from the DSS to maintain synchronization. If there is no DSS data as a result of an eclipse, the spin phase will be interpolated using a constant rate from the last DSS Sun pulse of one data set to the first Sun pulse of the next data set. It could also be extrapolated using the mean rate from the SpinKF.

With this approach, the resulting attitude solutions in the data gaps do account for coning, but not for nutation. Nutation damping results are covered in Section 5 below.

3.5. Attitude Estimation Results. The attitude estimation error has been found to be well under the required 0.1 deg (3σ) tolerance. Typical 3σ errors are roughly 0.05 deg about the Z-axis and 0.02 deg about X and Y. Sections 3.1, 3.2, and 3.3 presented the attitude and body rate estimation dependency on the knowledge of the inertia tensor and alignments of the four STS heads. Figure 1 below shows the degraded attitude solutions during the first couple orbits of MMS-3 with the pre-launch and non-calibrated star camera alignments. (This figure includes SpinKF residuals from two of the four STS heads, and so shows two distinct lines.) As seen in the figure, an offset relative to the mean body frame is observed for each of the STS head attitude residuals. The other MMS spacecraft and the other residuals components are qualitatively similar, although the separation between the two camera heads seen in Fig. 1 is one of the largest examples. The change in the density of points near the middle of the time span is due to a change in the telemetry downlink rate.

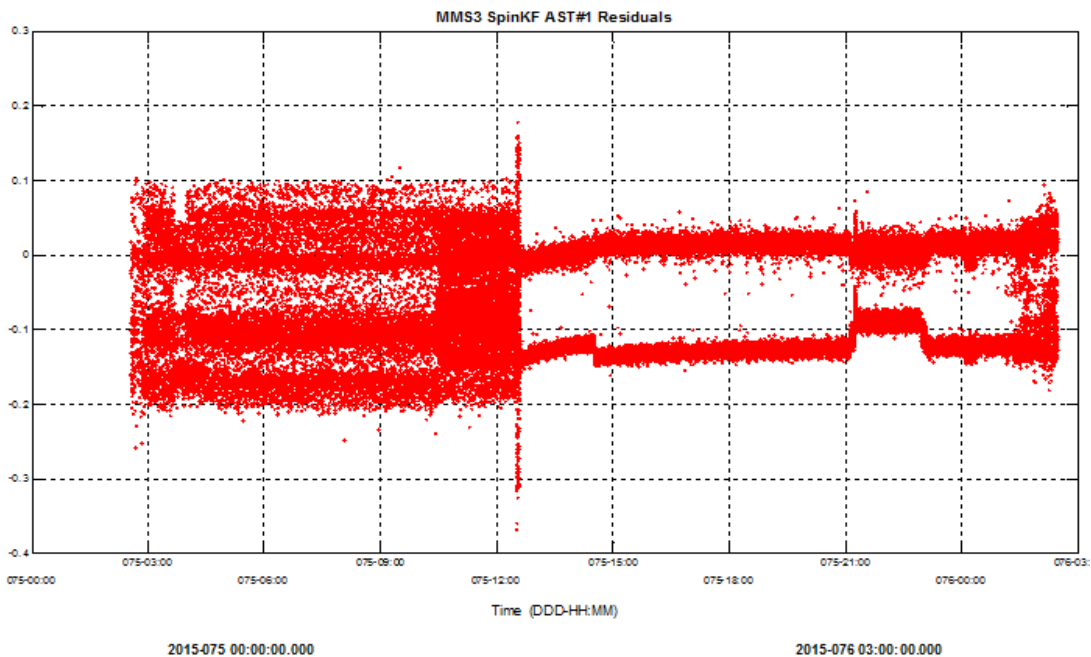


Fig. 1. Pre-calibration MMS3 attitude residuals for two of the STS heads (X-component, in degrees).

The offset between attitude solutions from each STS head disappears after the camera heads are properly aligned relative to the mean body frame (as defined using data from all four STS heads). This improvement in the residuals can be seen in Figs. 2 and 3; although, a small non-zero mean residual can still be seen in the attitude caused by uncertainty in the knowledge of the inertia tensor.

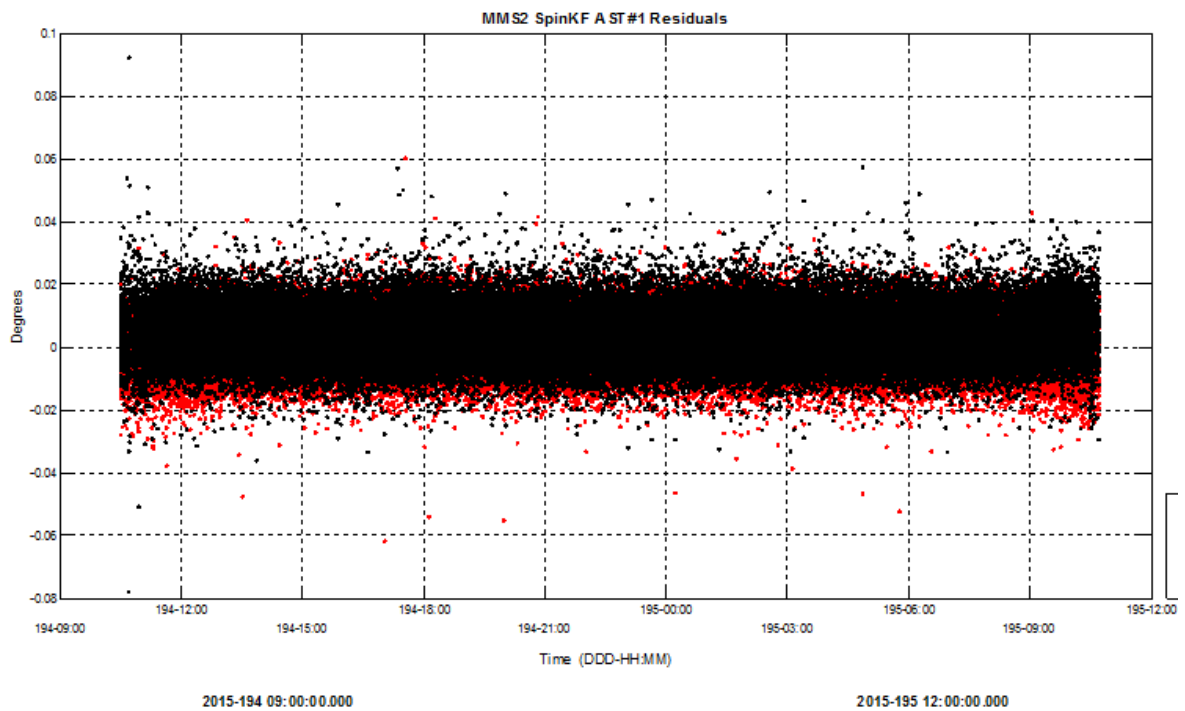


Fig. 2. Post-calibration MMS2 attitude residuals for two of the STS heads (X- and Y-components shown).

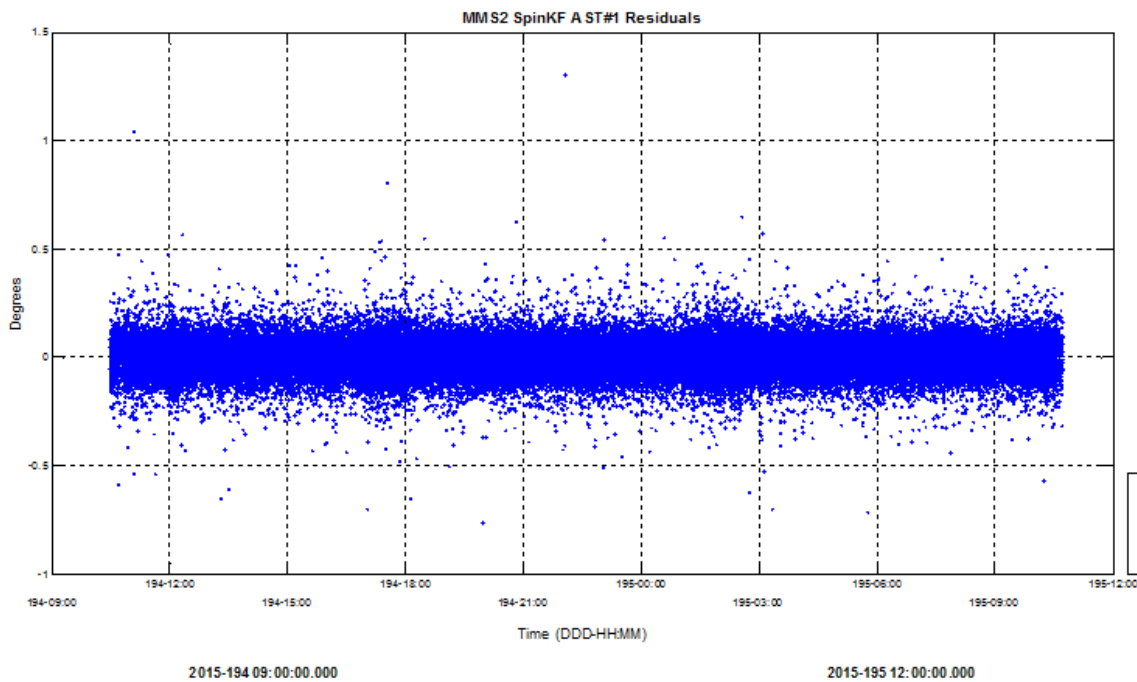


Fig. 3. Post-calibration MMS2 attitude residuals for two of the STS heads (Z-component shown).

Figure 4 shows evidence of how the attitude accuracy is affected by the MPA value. A maneuver occurs at 06:00 (where the spikes appear on the plot). The MPA was estimated using data after this maneuver, and thus, the MPA is appropriate for the mass and distribution of fuel after the burn. Then, the attitude was estimated for the entire time span using that post-burn MPA. As seen in Fig. 4, the attitude residuals shift closer to zero after the burn. This indicates the attitude estimate was slightly less accurate prior to the burn, using the slightly inaccurate post-burn MPA. This can be attributed to some amount of asymmetric fuel usage leading to a shift in the MPA. However, analysis done for smaller maneuvers suggests a nearly unnoticeable change in the major principal axis direction, as compared to larger maneuvers.

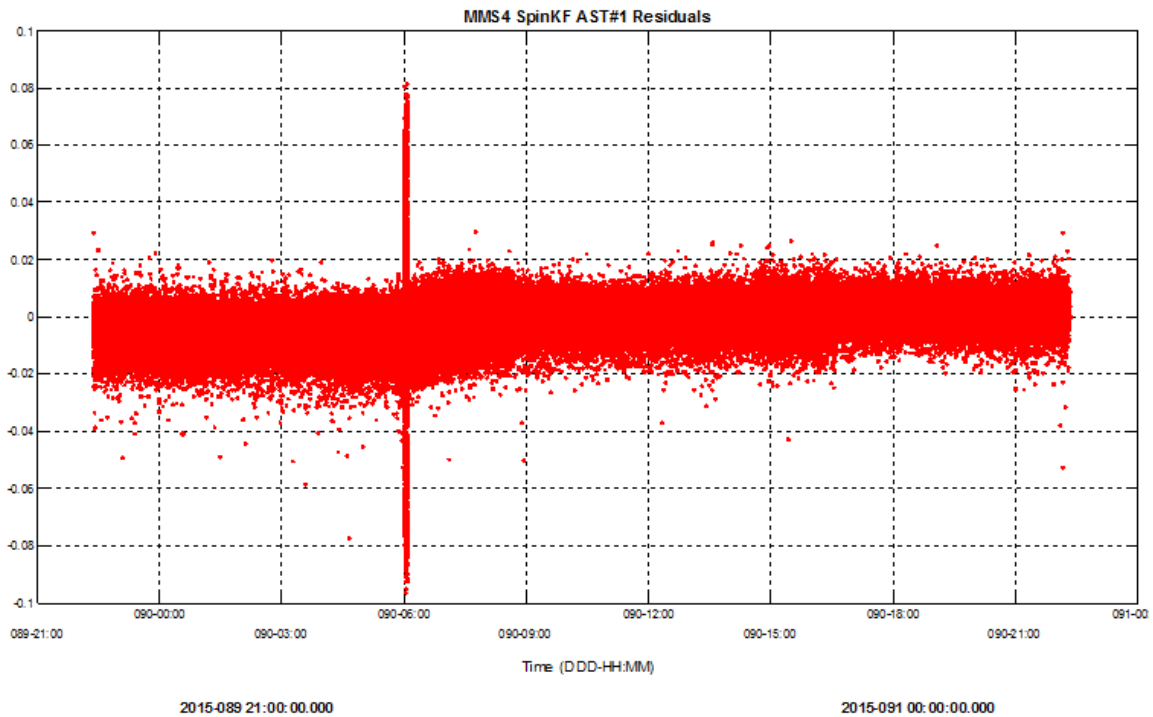


Fig. 4. MMS1 X-axis attitude residuals (degrees) for two of the STS heads, showing a shift during a maneuver.

Any small error in the MPA direction that causes non-zero mean attitude residuals also affects the body rate estimates. True misalignment of the MPA from the body Z-axis is the coning angle. On top of the true coning, there may be some small additional error in our knowledge of the MPA direction. Figure 5 shows how the spacecraft body rate estimates are biased (non-zero mean). It can be shown that this bias is in proportion to the coning angle times the spin rate. Figure 6 shows the spin rate about the body Z-axis. The deviation from the nearly constant rate seen near the end of the time span is due to the spacecraft passing through the shadow of the Earth. As the booms cool, they contract and the rate increases to conserve angular momentum.

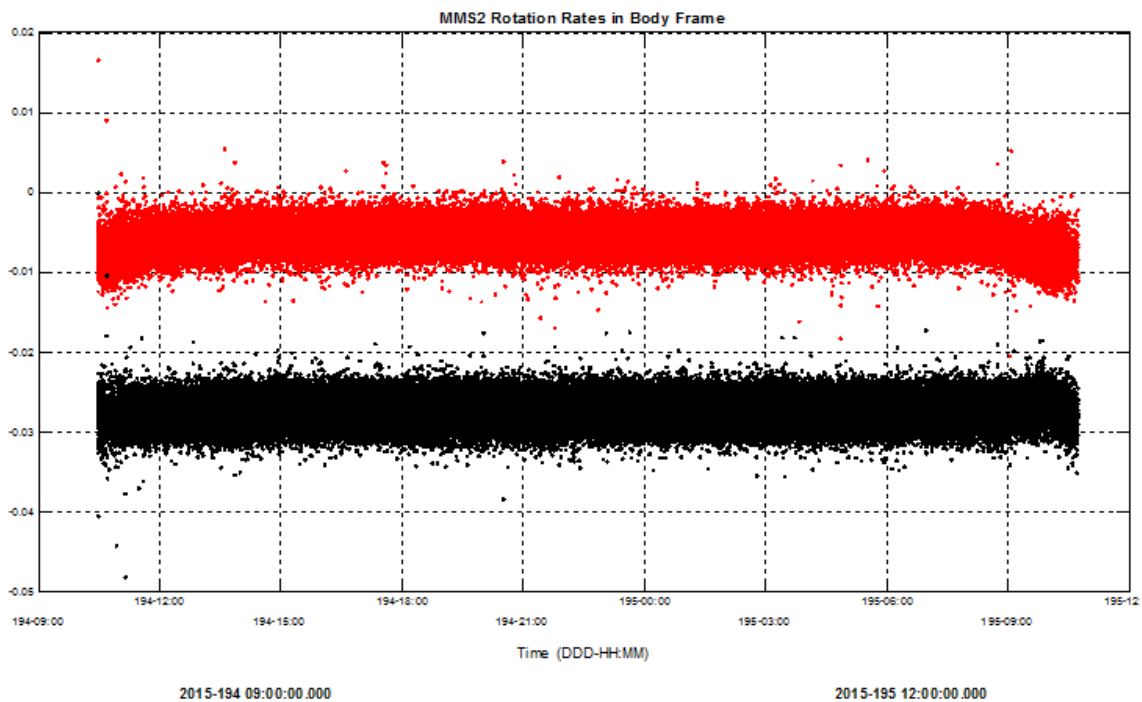


Fig. 5. MMS2 spacecraft body rates (red and black are the X- and Y-components, in deg/sec).




Fig. 6. MMS2 spacecraft body rates (Z-component).

Another error source that affects the attitude and body rate estimates similar to poor knowledge of the major principal axis direction and STS head misalignments is the tuning of the Kalman Filter, which is discussed in Section 6.

SpinKF statistics for the attitude, rate, and individual sensor residuals are saved in a trending database. This database serves two important purposes: it allows for early detection of signs of sensor degradation, and it shows how fast the angular momentum is precessing in the GCI frame due to environmental torques. The environmental torque model implemented in the AGS is described in Section 4. This model is used for “smart targeting” of the attitude maneuvers (also in Section 4) and for the attitude predictions needed for ground contact analysis.

Characterization of error sources with data collected from MMS launch until formation initialization shows the STS noise to be 18 to 25 arcsec about the X- and Y-axes (transverse) and 200 to 240 arcsec about the Z-axis (boresight). As mentioned earlier, some fuel asymmetry has been found between the four tanks when fuel is being depleted during maneuvers. This asymmetry is expected due to the shape the tanks take when fuel is used during maneuvers. This change in the MPA is deemed negligible for small maneuvers, though.

Analysis has demonstrated that the calibration of the MPA is not sensitive to whether or not the vibrations of the long, thin-wire booms have damped out, which was an early concern. (Both prelaunch tests with high-fidelity simulation data and post launch tests corroborate this result.) However, the results of inertia tensor calibration analyses consistently suggest a difference of 8-12 arcsec in the MPA direction between the calibrated and prelaunch mass properties. Multiple calibrations of the principal axis direction have been run to eliminate the possibility that the variation is simple statistical uncertainty in the results irrespective of fuel usage. This difference could then be attributed to multiple factors, as follow:

- Modeling error as there was no prelaunch measurement of the inertia tensor with all booms deployed.
- Fuel asymmetry and possibly variations in thermal deformations.
- Initial uncertainty of the mass properties after the tanks were fueled.

3.6. Inertia effects due to mass asymmetry. It is known that the ability of MMS to perform accurate orbit maneuvers depends in part on knowledge of the center of mass (CM). At present, the AGS CM model is based on prelaunch measurements with all booms stowed, an analytical model for deployed booms, and interpolation on the remaining mass of fuel in the tanks. Separate from this CM model, a recursive least-squares filter has been implemented for CM estimation using the Doppler shifts of the GPS signals. However, this CM filter has not yet been tested with flight data. A detailed description of the CM filter is given in [1].

It is expected that one of the possible errors in the CM and inertia tensor models is due to the distribution of fuel mass among the four tanks. The same amount of fuel should be used from each tank, but in practice there will always be some asymmetry. This section investigates how the CM and the inertia tensor are affected by the amount of mass asymmetry. The final results given below are that the CM is not changed significantly, but there is a 0.005 deg shift of the coning angle for each kilogram of fuel asymmetry (that is, for each kilogram of mass shifted from one fuel tank to the tank on the opposite side of the spacecraft). The actual location of that asymmetric mass is unknown, so this analysis should be interpreted as yielding only a rough approximation for the coning angle.

Begin by dividing the total rigid-body inertia tensor into contributions from the rigid core, the fuel, and the wire booms

$$I_{tot} = I_c + I_f + I_w . \quad (9)$$

(This breakdown in Eq. 9 is not really needed for the analysis, but it does help to see the relative sizes of the inertias.) Since only the effect of small changes in the mass are of interest here, the geometry is simplified by rounding the actual numbers, dropping the off-diagonal terms, and averaging the transverse components. Then, the rigid-body inertias for the MMS spacecraft are approximately

$$I_{tot} = \begin{bmatrix} 3240 & 0 & 0 \\ 0 & 3240 & 0 \\ 0 & 0 & 5460 \end{bmatrix} \text{kg-m}^2 , \quad (10)$$

where

$$I_c = \begin{bmatrix} 1250 & 0 & 0 \\ 0 & 1250 & 0 \\ 0 & 0 & 1490 \end{bmatrix} \text{kg-m}^2 , \quad (11)$$

$$I_f = \begin{bmatrix} 60 & 0 & 0 \\ 0 & 60 & 0 \\ 0 & 0 & 100 \end{bmatrix} \text{kg-m}^2 , \quad (12)$$

and

$$I_w = \begin{bmatrix} 1930 & 0 & 0 \\ 0 & 1930 & 0 \\ 0 & 0 & 3870 \end{bmatrix} \text{kg-m}^2 . \quad (13)$$

The rigid-body inertias, and not the effective inertias, are used in this analysis since coning is not a dynamical motion. (Applying this analysis with the effective inertia would yield a coning angle roughly twice as large.)

The mass of the fuel is approximately 400 kg (100 kg per tank) and the dry mass is 940 kg, so the total mass, M_{tot} , is 1340 kg. The center of mass of each fuel tank is roughly 0.5 m from the spacecraft center of mass in the XY-plane and near the spacecraft center of mass in the Z-direction.

To begin to think about fuel usage asymmetry, consider a quantity of fuel Δm displaced from one tank to the tank on the opposite side along the body Y-axis. (The actual tank orientation does not matter since cylindrical symmetry is assumed in Eqs. 10-13) The Y-distance of Δm from the center of mass is $\Delta y = 0.5$ m. The spacecraft center of mass will shift by $2 \Delta y \Delta m / M_{tot}$, which is

0.0007 meters for each kilogram shifted between tanks. (The factor of 2 appears because this imagines Δm taken from one tank and that same Δm added to the opposite tank.) The corresponding change in the moment of inertia about the Z-axis is given by the parallel-axis theorem as $M_{tot} (2\Delta y \Delta m / M_{tot})^2$, which is only $0.0007 \Delta m^2 \text{ kg}\cdot\text{m}^2$, with Δm given in kg. This shows that the change in the inertia tensor due to the change in the center of mass is negligible for this analysis.

With this same level of approximation, neither the Z-axis moment of inertia nor the transverse moment of inertia will change when the mass Δm is shifted, because the same amount of mass still exists with the same moment arms (just shifted from one side to the other). However, the off-diagonal inertia elements will change in proportion to Δm . If the mass is shifted along the Y-axis of the body, then the YZ-component of the inertia tensor will change by $-2\Delta m$ times the offset of that mass in the Y- and Z-directions from the center of mass. The Y-offset is taken to be the distance to the tank center, which is roughly 0.5 m. For this analysis, assume that the mass Δm is located near the center of one of the tank hemispherical caps. That would give Δm an offset in the Z-direction of roughly 0.3 m from the spacecraft center of mass. Thus, the change in the off-diagonal component of the inertia tensor caused by the mass asymmetry Δm is roughly

$$I_{yz} = -2\Delta m \Delta y \Delta z \approx -0.3\Delta m \text{ kg}\cdot\text{m}^2 \quad (\text{with } \Delta m \text{ in kg}). \quad (14)$$

The change in the off-diagonal inertia tensor given in Eq. 14 causes a change in the coning angle. This dependence is derived next under the assumption that that angle and Δm are small (they will be found to be proportional). The coning angle can be determined by diagonalizing the inertia tensor. The rotation angle needed to do the diagonalization is the rotation from the old MPA to the new MPA. The old MPA is the body Z-axis. Thus, this diagonalization rotation angle is the coning angle.

With the contribution from Eq. 14, the inertia tensor is

$$\tilde{I} = \begin{bmatrix} I_t & 0 & 0 \\ 0 & I_t & I_{yz} \\ 0 & I_{yz} & I_z \end{bmatrix}. \quad (15)$$

This can be diagonalized by a rotation about the X-axis by some angle ϑ ,

$$\begin{bmatrix} I_t & 0 & 0 \\ 0 & \tilde{I}_t & 0 \\ 0 & 0 & \tilde{I}_z \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c & s \\ 0 & -s & c \end{bmatrix}^T \begin{bmatrix} I_t & 0 & 0 \\ 0 & I_t & I_{yz} \\ 0 & I_{yz} & I_z \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & c & s \\ 0 & -s & c \end{bmatrix}, \quad (16)$$

where $s = \sin(\vartheta)$, $c = \cos(\vartheta)$, and the tildes indicate that the diagonal components may be affected. (They in fact change only in second-order, but that result is not needed for the current analysis.)

Considering just the (2,3)-element, one finds

$$0 = (c^2 - s^2)I_{yz} - s^2 I_z + s c I_t. \quad (17)$$

This can be expanded into a quadratic equation in s^2 (since $c^2 = 1 - s^2$). (Note that one solution is spurious since, for that case, Eq. 16 does not go back to Eq. 10 when I_{yz} goes to zero. There is an implicit divide by zero in that limit for that solution.) However, under the assumption that the angle ϑ is small, Eq. 17 can easily be solved in first-order to find

$$\vartheta \approx s \approx -\frac{I_{yz}}{I_t}. \quad (18)$$

Combining Eqs. 10, 14, and 18, one finds $\vartheta \approx 0.3 \Delta m / 3240$ (Δm in kg, ϑ in radians). This yields a coning angle of 0.005 deg for each kilogram of mass asymmetry.

4. Environmental Torque Model

The need for attitude maintenance maneuvers is driven by a combination of attitude drift and the seasonal motion of the Sun in the inertial reference frame. The target spin axis for science support is limited to a small box, offset from the ecliptic pole by 3.5 deg and having dimensions of approximately 2.5 deg \times 2.5 deg. The AGS must predict when the spin axis will drift to the edge of this box and then plan a slew to the opposite side of the box in order to maximize the time between attitude maneuvers. (This is the “smart targeting” algorithm in the AGS, which must account for the Sun and target box geometry and the seasonally changing environmental torques.)

The dominant environmental torque driving the attitude drift is caused by gravity-gradient forces. The torques from atmospheric drag and solar pressure are smaller by a few orders-of-magnitude. Very roughly, the order-of-magnitude of the gravity-gradient torque is 10^{-4} N-m, the solar pressure torque is less than 10^{-6} N-m, and the aerodynamic drag torque is less than 10^{-7} N-m.

An expression for the spin-averaged and orbit-averaged gravity-gradient torque is

$$\boldsymbol{\tau}_{gg} = \frac{3}{2} \frac{\mu}{a^3(1-e^2)^{3/2}} (I_z - I_t) (\hat{\mathbf{Z}} \cdot \hat{\mathbf{h}}) (\hat{\mathbf{Z}} \times \hat{\mathbf{h}}), \quad (19)$$

as given in [10] (with corrected factor of 1/2), where μ is the Earth’s gravitational constant, a is the semimajor axis, e is the eccentricity, I_z is the principal moment of inertia, I_t is the transverse moment of inertia (taken as the mean of the x - and y -inertias), \mathbf{Z} is the body Z -axis, \mathbf{h} is the orbit normal vector, and the caret indicates unit vector.

Application of Eq. 19 to the MMS angular momentum vector yields predictions for the spin axis precession. These predictions have been compared to the observed drift of the spin axis at various stages of the MMS mission. Table 1 shows a sample of these results for MMS1. The first row shows results for the precession after deployment of the magnetometer boom, but prior to deployment of the long, thin-wire, spin-plane double-probe (SDP) booms. The next row shows results for times during the multi-step SDP deployment, and the last row shows results after SDP deployment was complete. The SDP is significant because its deployment increases the value of $I_z - I_t$ by 60 percent. (However, it is worth noting that the *ratio* of these inertias does not change significantly with SDP deployment, and it is the ratio that most affects the rate of precession.)

Table 1. MMS1 predicted and observed angular momentum precession averaged over several days.

Deployment Status	Predicted Precession per Day (deg)	Observed Precession per Day (deg)	Error in Precession per Day (deg)
Pre-SDP Deployment	0.0856	0.1052	0.0196
During SDP Deployment	0.0565	0.1200	0.0056
Post-SDP Deployment	0.0514	0.0513	0.0015

The last column in Table 1 shows that the predictions of the angular momentum precession improved as the SDP was deployed. This improvement can probably be attributed to a more accurate knowledge of the actual inertia tensor for the deployed configuration. Note that the errors shown in the last column are not simply the difference between the second and third columns. The latter are precession magnitudes averaged over several days, while the precession errors are computed as the average of the actual angle differences between angular momentum vectors.

Figure 7 presents the accumulated errors in the predicted precession angles for MMS2. This plot shows a long time span when no maneuvers were performed. The results are similar or better for the other three observatories. These errors are smaller than the prelaunch error estimates by an order of magnitude due to better than anticipated knowledge of the inertia tensor values.

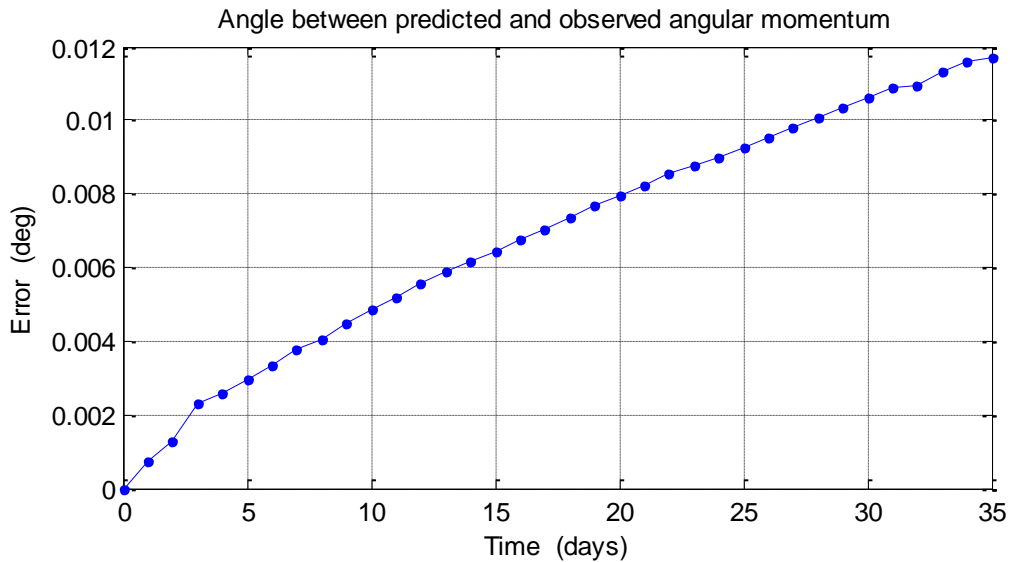


Fig. 7. Angle between predicted and observed angular momentum vectors for MMS2.

5. Nutation and Vibration Damping

Every thruster burn is expected to excite nutation of the rigid core of the spacecraft and vibration of the flexible appendages. Some level of excitation is also expected every orbit due to the rapidly changing gravity-gradient torque during perigee passes. In addition, when the spacecraft enters or exits Earth's shadow, the booms change temperature and contract or expand slightly, causing a spin rate change that excites a particular mode of vibration.

Prior to launch, no detailed damping model was available to predict damping times for nutation and vibration. Since launch, it has been possible to measure the damping times at various stages of boom deployment. After damping is essentially complete, the angular momentum becomes aligned with the major principal axis (MPA) of the spacecraft, and will remain constant in the absence of external torques or thruster pulses. While damping is ongoing, the angular momentum will still be constant in inertial space, but the MPA vector will precess about that direction, spiraling in with some characteristic decay time. The mean MPA direction will very nearly equal the true angular momentum direction when averaged over many nutation or vibration periods. Hence, the decay time can be determined by computing the offset angle between the mean and instantaneous MPA vectors. If this offset angle decays exponentially, then a linear fit to its logarithm will yield the decay constant.

Figure 8 shows an example of the MPA offset angles for MMS3 after deployment of the magnetometer booms. No other booms are deployed at this early stage of the mission, so this figure shows nutation of the nearly rigid spacecraft. The actual vibrations of the magnetometer booms themselves are higher in frequency than the spin rate and they decay on a much shorter time scale. Figure 9 presents the logarithm of the MPA offset angle and its linear fit (red line). From this curve, the half-amplitude decay time is found to be 0.46 hours. Similar values are found for all four MMS spacecraft for the rigid-body nutation damping prior to deployment of the wire SDP booms.

Table 2 gives the half-amplitude decay times at three different stages in the boom deployment history. The first row shows the nutation damping after the magnetometer boom deployment, as discussed above. The second row shows the damping after a spin adjust maneuver at a time after the full deployment of the SDP radial wire booms. Two things can be noted here. The damping time is longer than for the magnetometer booms by factors ranging from 10 to 30, and the values for the four spacecraft are all rather different. The longer damping times are due to the larger moments of inertia after the SDP deployment and the presence of several low frequency modes of SDP oscillation. The differences in damping times for the four spacecraft have not yet been satisfactorily explained.

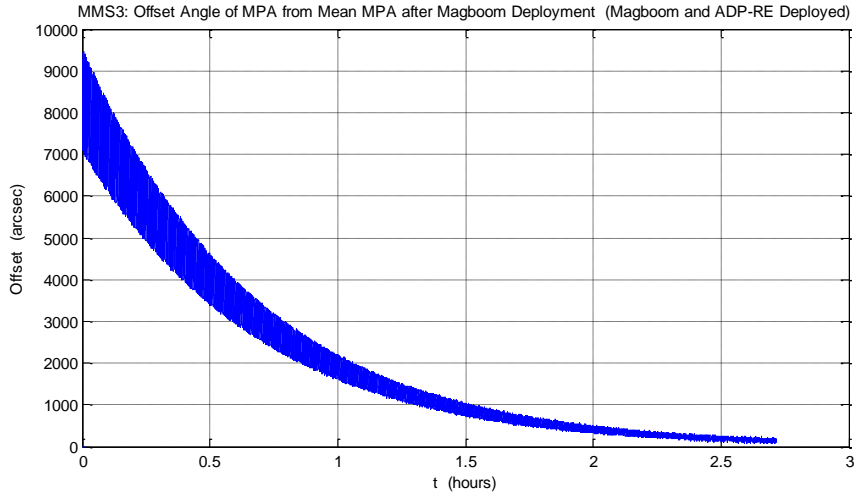


Fig. 8. Angle (arcsec) between mean and instantaneous direction of the major principal axis for MMS3 after deployment of the magnetometer booms.

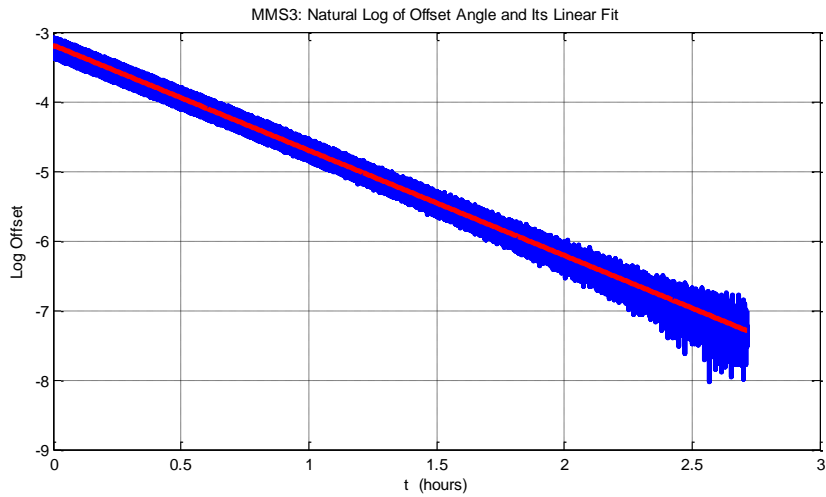


Fig. 9. Log of the offset angle after deployment of the magnetometer booms. Red line shows linear fit.

The third row of Table 2 shows the damping times after all booms are deployed: magnetometer booms, radial SDP wire booms, and the axial double-probe (ADP) booms. It is found that the times given in Table 2 can vary by as much as 25 percent depending on how long it has been since the last burn. Multiple vibrational modes can be excited, each having its own decay time. Given this variability and uncertainty in the damping times, the times with all booms deployed are not too different from the second row, with the ADP stowed. The ADP has sufficient rigidity so that its dominant vibrational modes are much higher in frequency than the SDP vibrations, and they are expected to damp much faster. For this reason, the ADP does not contribute very much to the overall damping time.

Table 2. Half-amplitude decay times for MMS nutation and vibration after various deployments.

Deployment Status	MMS1 Decay Time (hrs)	MMS2 Decay Time (hrs)	MMS3 Decay Time (hrs)	MMS4 Decay Time (hrs)
Mag Boom Deployed (SDP and ADP Stowed)	0.38	0.41	0.46	0.44
Mag Boom and SDP Deployed (ADP Stowed)	10.0	12.8	4.3	8.2
All Deployed	6.5	11.5	5.4	8.3

Between burns there is no active damping; only the internal friction in the fuel tanks and in the booms themselves contribute to the decay. The long decay times and the energy input at each perigee pass imply that there will always be some low level of vibration; however, it is found that the amplitude of this background boom excitation is less than 0.01 deg and cannot be distinguished from noise and other error sources in the attitude estimator.

6. Kalman Filter Tuning

The AGS delivers a daily definitive attitude and rate product. This product is generated using a Kalman filter designed for spinning spacecraft. This filter, called the SpinKF, is described in detail in Refs. 5 and 6. A high-level description is given in [1].

A few years prior to launch, the SpinKF was tuned by adjusting the two process noise parameters to optimize filter performance using data from a simple rigid-body MMS simulator. These noise parameters represent an angle random walk and a rate random walk. Data became available later from the GSFC MMS constellation high-fidelity simulator (CHIFI). This simulator includes detailed sensor modeling with realistic noise and digitization. The dynamics model includes options for nutation, coning, and vibrations of the wire SDP booms and other appendages.

To generate CHIFI data for testing the AGS, a small attitude maneuver burn would be performed to excite boom vibrations at the start of each run. Several test cases were created to verify the AGS could satisfy mission requirements for attitudes and body rate estimates. With this data, it was found that the standard deviations of the star camera residuals from the attitude estimation could be significantly reduced by decreasing the rate random walk process noise by an order of magnitude.

The values for the angle random walk noise source strengths used at launch were 3×10^{-9} rad²/sec about the spin axis and 1×10^{-9} rad²/sec about the transverse axes. The rate random walk noise strengths were 1×10^{-9} rad²/sec³ on all three axes. The sensor noise parameters for the star cameras were set to their specification values, 20 arcsec transverse to the camera boresight and 67 arcsec (1σ) around the boresight.

After launch, it was found that the star camera noise transverse to the boresight is near the specified 20 arcsec, but noise around the boresight is approximately three times larger than expected.

A series of tests were run using flight data with the star camera sensor noise parameter increased to 200 arcsec about the boresight. This increase improved the star camera Z-axis residuals from the Kalman filter by 2 to 10 percent, depending on the choice of process noise parameters.

The process noise parameters were separately varied over several orders of magnitude. The SpinKF filter was executed for a two-hour time span for each case, and the star camera residuals statistics were computed. The time span is one day after an attitude maneuver so boom vibrations have not yet fully damped out. Figure 10 shows a surface plot of the standard deviations of the star camera Z-axis residuals. The X- and Y-axis residuals are qualitatively similar.




Fig. 10. Standard deviations of star camera Z-axis filter residuals for varying process noise parameters.

Filter tuning requires a trade-off between achieving smooth solutions where the sensor noise is removed and not reducing the state covariance so much that the filter responds unacceptably slowly to any actual changes. Too small a covariance can also lead to problems in filter recovery if the filter estimate is pushed far from the true state due to anomalous or outlying data. This trade-off means choosing small process noise for smooth solutions, but not too small so the filter remains responsive.

Figure 10 shows that a choice of 1×10^{-9} rad²/sec angle random walk noise and 1×10^{-10} rad²/sec³ rate random walk noise yields small filter residual errors. This new tuning leaves the angle random walk unchanged and reduces the rate random walk by an additional factor of 10. Further reduction of these parameters decreases the residuals only negligibly.

Figure 11 compares the raw star camera measurements with the corresponding filter Euler angle estimates for a one minute time span. The oscillations are due to boom vibrations. It is clear that the random noise seen in the raw star camera measurements on the left in Fig. 11a is largely removed in the filter estimates shown on the right in Fig. 11b.

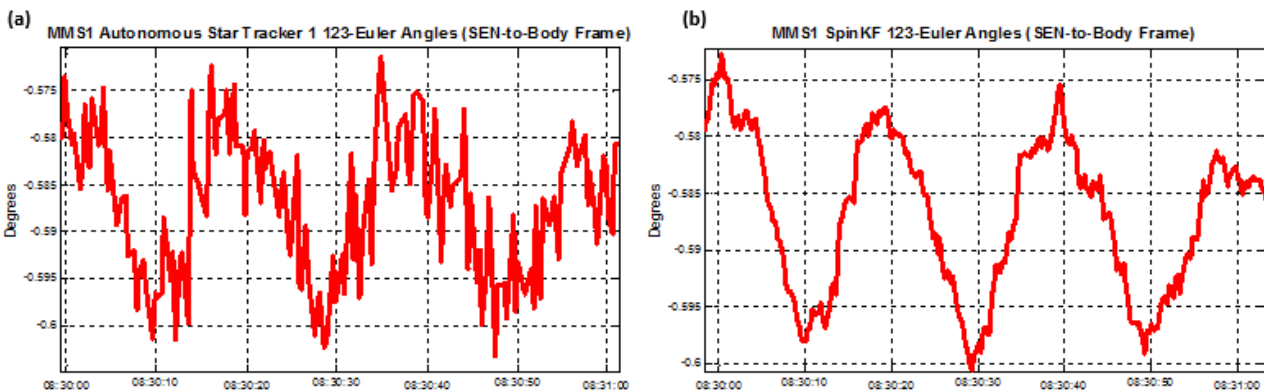


Fig. 11. (a) Star camera X-axis measurements. (b) Filter estimate for the same Euler angle component.

Further reduction of the process noise could remove the last of the random noise seen in Fig. 11b, but the improvement in the residuals would be very small and it would make the filter less responsive with no significant benefit. Furthermore, the definitive attitudes and rates delivered to the Science Operations Center (SOC) since launch have been well within specifications using the prior tuning. However, this new tuning should provide even smoother solutions for the pointing directions, rates, and spin phase that are delivered to the SOC daily.

7. Shadow Period Attitude Optimization

Early in the mission, the four MMS spacecraft experienced a “shadow period” lasting roughly two months, during which the spacecraft received relatively low levels of sunlight and needed therefore to remain in a low-power state. To support a low-power state, no maneuvers could take place during the shadow period. Therefore, it was desired that the last available attitude maneuver opportunity before the shadow period be used to put the spacecraft in the best possible attitude. The best target attitude for this purpose is determined by minimizing the angle between the spin axis direction and the center of the mission target box over the duration of the shadow season (see Section 4 for a detailed description of the target box).

Other than during the shadow period, attitude maneuvers are performed periodically to keep the spin axis within the target box. Remaining close to this attitude is important for spacecraft power and thermal reasons, and for science instrument health and safety.

The Spinner Attitude Prediction (SAP) utility of the MMS AGS was used as the simulation environment for this study. Given an initial attitude, SAP is able to predict the spin axis direction at apogees and perigees for a selected time span by using a model for the gravity gradient torque, the dominant environmental torque on the spacecraft (see Section 4 of this paper and also [10]).

To solve for the optimal initial attitude for the shadow period, a cost function was chosen that gives preference to spin axis directions nearer to the center of the target box. That is,

$$C = [\sum_{i=1}^N (\cos^{-1}(\mathbf{v}_i \cdot \mathbf{w}_i))^2]^{1/2} \quad (20)$$

where N is the number of time steps, \mathbf{v}_i are the spin axis direction unit vectors predicted by the SAP utility at each time step, and \mathbf{w}_i are the corresponding unit vectors in the direction to the center of the target box.

Equation 20 is particularly useful because it reduces the problem to the minimization of a scalar function. The built-in MATLAB function `fminsearch` was used to find the attitude at the start of the shadow season that minimizes the function for the selected time period.

One can visualize the state space of the cost function from Figs. 12a and 12b, which actually show the complement of the cost function (the “value” function). In these figures, maxima correspond to the most desirable attitude targets. The two horizontal axes are the right ascension and declination of the initial spin axis direction, and the vertical axis is the value assigned to that initial spin axis. The optimal attitude returned by `fminsearch` is shown as the green point.

The value function was observed to be unimodal, which makes it easier to search the state space for the optimal value. Figure 12b is particularly interesting because its scope was chosen to be small enough to expose the resolution limitation of the model, shown by the flat quadrilateral sections. This granularity appears to be due to the discrete time steps used in SAP when propagating the initial attitude.

A small error due to the optimization search function is also evident in Fig. 12b. The green marker indicates the solution reached by `fminsearch`, which evidently does not coincide perfectly with the maximum. This error is on the order of a few ten-thousandths of a degree, which is completely negligible when the accuracy of the model or the requirements of the mission are taken into account.

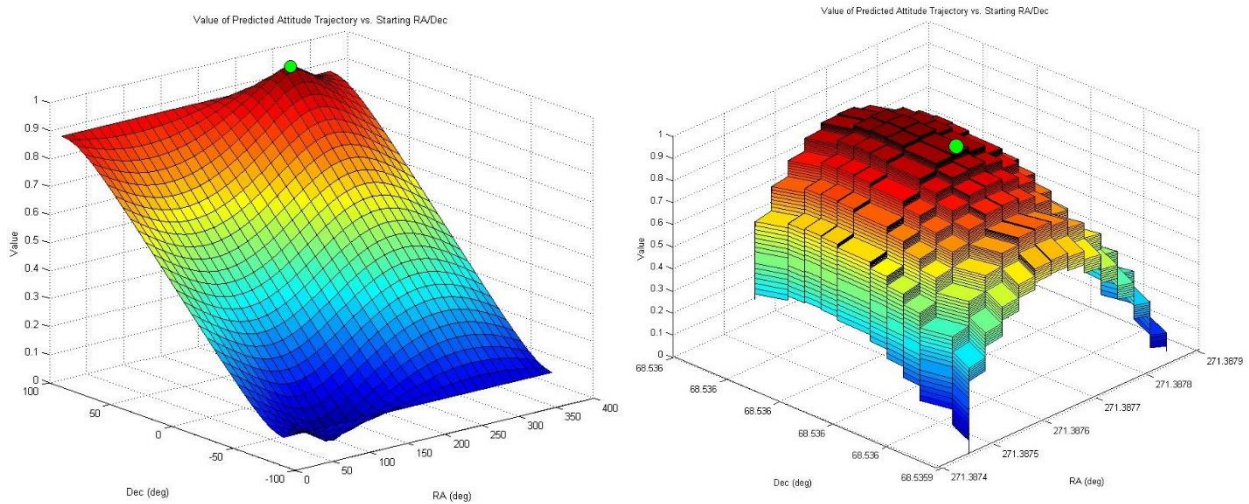


Fig. 12. a) Surface plot of the value function; b) detail of the value function near its maximum.

The result of the shadow period attitude optimization can be seen schematically in Fig. 13, which shows a typical predicted attitude trajectory over the early-mission shadow period after maneuvering to the optimal attitude. Later in the mission, there may be similar shadow seasons where again it will be necessary for MMS to operate in low-power mode. This optimization method will then again be useful for targeting the optimal initial attitude.

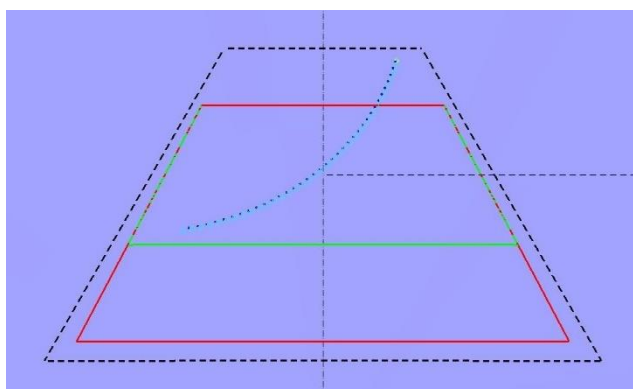


Fig. 13. Shadow period spin axis trajectory (blue circles) resulting from an optimized initial attitude. The black dashed line signifies the target box.

8. Earth Albedo Sensor Interference

After the negative Z-axis ADP booms were deployed, unexpected periodic interference was observed in the STS data for all four spacecraft. The interference was manifested as a significantly higher incidence of flagged (bad) data points coming from the STs.

The hypothesis is that the observed interference is due to glint or diffuse reflection from the negative Z-axis ADP boom entering the STS Sun shades and interfering with the star cameras. It is known that the ADP boom receives illumination from both sunlight and Earth albedo, but this analysis focuses only on the Earth albedo contribution since that part of the illumination is time-dependent (depending on the spacecraft position relative to Earth). The three primary clues leading to the hypothesis are: 1) the interference was not observed until the negative Z-axis ADP boom was deployed, 2) the interference occurs near every perigee (when the effects of Earth albedo are at a maximum), and 3) the interference is correlated with spin phase, which suggests that it may be related to the geometry of the spacecraft.

Prior to any analysis to find a method of modeling the observed sensor interference, a method of identifying intervals of interference was required. An algorithm was devised that could automatically identify interference by looking at the elapsed time between good (non-flagged) data points in the raw telemetry. Figure 14 shows an illustration of how the algorithm perceives this data. An automatically detected interference interval is annotated by the red boundaries.

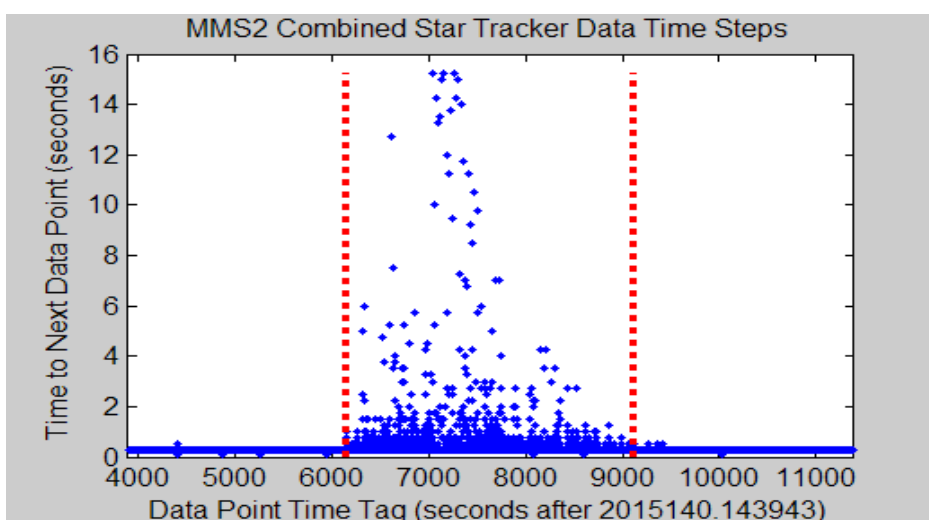


Fig. 14. Time between non-flagged STS data points. Red lines indicate boundaries of interference interval.

Since a rigorous analysis of the spacecraft geometry is impractical, it was decided to use artificial intelligence techniques to determine empirically what conditions cause the interference. This

approach required significantly less time investment, and was especially convenient because the automatic interference detection algorithm could be used to implement an unsupervised machine learning routine.

The problem was approached as a classification problem where the input space is comprised of parameters describing the spacecraft state, and each point in the input space falls into one of two classes: A) interference, or B) no interference. In order to follow through with this approach, sensible parameters had to be chosen and calculated. Since it is hypothesized that the cause of the interference is dominated by Earth albedo effects, the following parameters were considered:

- X_1 : Angle from the body negative Z-axis to the direction of nearest lit point on Earth
- X_2 : Total visible angular area (from the spacecraft’s perspective) occupied by the lit part of the Earth
- X_3 : Fraction of the visible Earth disk (from the spacecraft’s perspective) that is lit

In the parameters described above, the “lit” Earth refers to the portion of the Earth’s surface that is in direct sunlight. The body negative Z-axis was chosen for the first parameter because all four STS point generally in that direction (each camera boresight is roughly 10 degrees offset from the body negative Z-axis).

One popular and successful artificial intelligence framework is that of the Support Vector Machine (SVM). Quoting from Ref. 11, “SVMs construct a maximum margin separator – a decision boundary with the largest possible distance to example points.” The decision boundary is a hyperplane that lies in the input space, separating the parameters by their classification (a “separating hyperplane”). Details giving the algorithm derivation and implementation may be found in [11]. The separating hyperplane can be described by a normal vector \vec{w} and a scalar *bias*. To classify a new point \vec{x}_{new} , the decision table shown in Table 3 is evaluated.

Table 3. Decision table describing model for predicting Earth albedo interference in the star cameras.

Condition	Classification
$\vec{x}_{new} \cdot \vec{w} - bias > 0$	Interference
$\vec{x}_{new} \cdot \vec{w} - bias < 0$	No Interference
$\vec{x}_{new} \cdot \vec{w} - bias = 0$	Undefined (point lies on decision boundary)

The challenge was to select an input space (some combination of parameters X_1 , X_2 , and X_3) that is sufficiently separable by a linear decision boundary (a hyperplane), and to have an SVM find the optimal decision boundary in that input space for a given set of training data. Eleven orbits of telemetry and definitive attitude data was obtained to generate a training set, and each non-trivial combination of the three parameters was considered as a candidate input space. For each candidate input space, an SVM was trained on the dataset represented in that input space, and a statistical study was done on the results to determine whether a linear decision boundary is valid.

The parameter combination $\{X_1, X_2\}$ was found to be a good input space for modeling the observed sensor interference with a linear decision boundary. Figure 15 shows the training dataset represented in this input space. Points classified as “interference” are shown in red, and points classified as “no interference” are shown in green. Though not shown in the figure, since this input space is 2-dimensional, the decision boundary would be a 2-dimensional hyperplane, more commonly referred to as a line. Note that classification of the training set was performed automatically by the interference detection algorithm described earlier, making this SVM training an application of unsupervised machine learning.

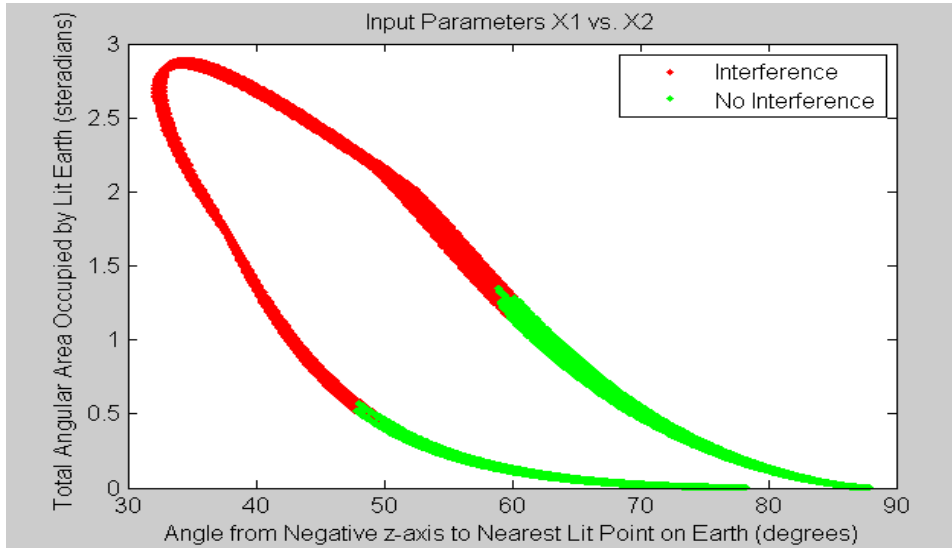


Fig. 15. Training data represented in the selected input space $\{X_1, X_2\}$, with classifications shown.

The results of SVM training based on the eleven-orbit dataset are shown in Table 4, where \vec{w} and *bias* describe the decision boundary as a separating hyperplane in input space $\{X_1, X_2\}$. This definition of \vec{w} and *bias* was used in conjunction with Table 3 to model interference. The model has been tested extensively using the available star camera data and has been shown to successfully predict interference due to Earth albedo illumination of the negative Z-axis ADP boom.

Table 4. Results of SVM training based on eleven orbits of data, showing weight vector \vec{w} and *bias*.

w_1 <i>radians</i> ⁻¹	w_2 <i>steradians</i> ⁻¹	<i>bias</i> <i>dimensionless</i>
-49.766	11.249	-38.522

The SVM framework is capable of solving classification problems for much more complex input spaces if a nonlinear kernel function is chosen. This study may be revisited later in the mission when a wider variety of training data is available.

9. Summary and Conclusions

The existing functionalities and new utilities implemented in the AGS have proven to be essential for MMS mission support in many ways. The support has included validation of the onboard spacecraft attitude and rate estimates, calibration of the star camera head alignments, calibration of the direction of the major principal axis of inertia, and attitude maneuver planning of “smart targets” to stay within the mission science target box. The AGS capabilities have enabled the proper identification of error sources for improving the attitude and body rate solutions, including thermal variations and sensor interferences that affect the attitude solutions.

The AGS was also used to perform some special analyses to support the mission. Namely, a long-term planning tool was developed in order to cope with the long shadow period near the beginning of the mission, and a method of sensor interference prediction was developed after observing unexpected sensor noise in the STS. It was found that the existing functionalities of the AGS and its library routines could be readily extended to solve both of these special-case problems.

Although there is high dependency of the attitude and rate solutions on the accuracy of the STS alignments and the knowledge of the inertia tensor, the SpinKF has performed very well to meet the MMS estimation tolerances. The dependency arises first from a need for self-consistent observations. These are provided by the relative STS alignment calibrations. Secondly, the use of Euler’s equations

of motion to propagate the state between observations requires a good inertia tensor. The MPA calibrations performed after every burn ensure the AGS always is using the best available inertia tensor and coning angle estimate. With these calibrations, the definitive attitude reports generated using the SpinKF have been delivered to the scientists daily, covering every orbit perigee-to-perigee, with the best available attitude accuracy.

It was known from early tests that there is coupling between the estimates for the STS alignments and the MPA direction. This was dealt with by iterating between these two calibrations. Once a set of STS alignments were found that made the four STS camera head measurements consistent, these alignments have been used both onboard and in the AGS. These alignments, in effect, define the MMS body frame, and the MPA calibrations have been performed relative to this frame ever since the alignments were uplinked.

The AGS requirement to support maneuver planning drove the need to predict the precession of the angular momentum for up to several weeks at a time. The AGS utility to predict the attitude relies on good knowledge of the inertia tensor along with an environmental torque model. This model has been shown to predict the spin axis direction very accurately, with an error of only 0.01 deg after a one-month propagation.

In summary, the AGS suite of utilities has been crucial to satisfying the MMS mission requirements, and has been used to provide good ground attitude determination and support to the mission during spacecraft commissioning and nominal operations.

References

- [1] Sedlak, J. E., Superfin, E. A., and Raymond, J. C., "Magnetospheric Multiscale (MMS) Mission Attitude Ground System Design," *22nd Int. Symposium on Spaceflight Dynamics*, INPE, São José dos Campos, SP, Brazil, Feb., 2011.
- [2] Langston, J., et al., "A Multimission Three-Axis Stabilized Spacecraft Flight Dynamics Ground Support System," *1992 Flight Mechanics/Estimation Theory Symposium*, NASA Conference Publications CP-3186, NASA/GSFC, Greenbelt, MD, May 1992.
- [3] Hashmall, J. A. and Sedlak, J. E., "New Attitude Sensor Alignment Calibration Algorithms," *53rd Int. Astronautical Congress*, IAC-02-A.4.07, IAF, Houston, TX, Oct. 2002.
- [4] Shuster, M. D. and Oh, S. D., "Three-Axis Attitude Determination from Vector Observations," *J. Guidance and Control*, Vol. 4, No. 1, Jan.-Feb. 1981, pp. 70-77.
- [5] Markley, F. L. and Sedlak, J. E., "Kalman Filter for Spinning Spacecraft Attitude Estimation," *Journal of Guidance, Control, and Dynamics*, Vol. 31, No. 6, p. 1750, Nov.-Dec. 2008.
- [6] Sedlak, J. E., "Spinning Spacecraft Attitude Estimation Using Markley Variables: Filter Implementation and Results," *2005 Flight Mechanics Symposium*, NASA Conference Publications CP-2005-212789, NASA/GSFC, Greenbelt, MD, Oct. 2005.
- [7] Lefferts, E. J., Markley, F. L., and Shuster, M. D., "Kalman Filtering for Spacecraft Attitude Estimation," *Journal of Guidance, Control, and Dynamics*, Vol. 5, No. 5, pp. 417-429, 1982.
- [8] Markley, F. L., "New Dynamic Variables for Momentum-Bias Spacecraft," *The Journal of the Astronautical Sciences*, Vol. 41, No. 4, pp. 557-567, 1993.
- [9] Shuster, M. D., "A Survey of Attitude Representations," *Journal of the Astronautical Sciences*, Vol. 41, No. 4, Oct.-Dec., 1993, pp. 439-517.
- [10] Wertz, J. R., (ed.), *Spacecraft Attitude Determination and Control*, D. Reidel Publishing Company, Dordrecht, The Netherlands, 1978.
- [11] Russell, S. J., and Norvig, P., *Artificial Intelligence: A Modern Approach*, 3rd Edition, Dec. 2009, pp. 744-747.