

## Chapter 6

### Attitude Determination

When developing a simulation, empirical data from the modeled system is useful for comparison. However, in the case of an attitude simulation, the raw observations from the satellite must be converted to an estimate of the satellite attitude before any simulation-to-empirical-data comparisons may be performed. This chapter develops an attitude determination filter which may be applied to any PMAC satellite. After the filter is defined, it is applied to on-orbit measurements to determine the attitude of a PMAC satellite. This provides the empirical data needed for comparison with the attitude simulation (developed in Chapter 8).

The filter is tuned using input from the CSSWE CubeSat (see Chapter 4) but is applicable to any satellite with a PMAC system when rate gyro data are not available. CSSWE is designed to perform attitude determination via post-processing using raw measurements transmitted to its ground station. However, given proper ground-based measurements before launch, the attitude determination developed in this chapter could occur in real-time on orbit.

#### 6.1 Filter Design

The Extended Kalman Filter (EKF) is an established method of attitude determination [46] [16]. In general, an EKF is useful for estimating the state and covariance of a non-linear, discrete-time process. What follows is an overview of the general EKF. First, assume the state propagation for a given process is governed by the non-linear stochastic difference equation

$$\mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}, \mathbf{w}_{k-1}) \quad (6.1)$$

with measurements

$$\mathbf{y}_k = \mathbf{h}(\mathbf{x}_k, \mathbf{v}_k) \quad (6.2)$$

where  $\mathbf{x}$  is the true state vector,  $\mathbf{u}$  is the true control input,  $\mathbf{w}$  is the true process noise,  $\mathbf{v}$  is the true measurement noise, and  $k$  is the step number. Note that each measurement is an inseparable combination of the current state and the measurement noise. The true process noise and measurement noise cannot be observed directly for each time step; instead the Kalman filter assumes all noise is Gaussian, independent, and zero-mean. Because Equation 6.1 represents a non-linear process it is difficult to propagate directly. The EKF state propagation procedure is shown below. Note that for the EKF matrices, the normal convention within this dissertation of matrices being bracketed is ignored.

$$\hat{\mathbf{x}}_k^- = \mathbf{f}(\hat{\mathbf{x}}_{k-1}^+, \mathbf{u}_{k-1}, 0) \quad (6.3)$$

$$P_k^- = F_k P_{k-1}^+ F_k^T + G_k Q_{k-1} G_k^T \quad (6.4)$$

$$K_k = P_k^- H_k^T (H_k P_k^- H_k + J_k R_k J_k^T)^{-1} \quad (6.5)$$

$$\hat{\mathbf{x}}_k^+ = \hat{\mathbf{x}}_k^- + K_k (\mathbf{y}_k - \mathbf{h}(\hat{\mathbf{x}}_k^-, 0)) \quad (6.6)$$

$$P_k^+ = (I - K_k H_k) P_k^- \quad (6.7)$$

where a hatted variable ( $\hat{\cdot}$ ) represents an estimate, a superscript minus ( $-$ ) or plus ( $+$ ) represents the *a priori* or *a posteriori* estimate before or after the measurement update,  $P$  is the state covariance

matrix,  $Q$  is the process noise covariance matrix,  $R$  is the measurement covariance matrix (not to be confused with the rotation matrix  $[R]$ ),  $F$  is the state transition matrix,  $G$  is the process noise gain matrix,  $H$  is the measurement sensitivity matrix, and  $J$  is the measurement noise gain matrix.

The Jacobian matrices are defined as follows:

$$F_k \equiv \frac{\partial \mathbf{f}}{\partial \mathbf{x}} (\hat{\mathbf{x}}_{k-1}^+, \mathbf{u}_{k-1}, 0) \quad (6.8)$$

$$G_k \equiv \frac{\partial \mathbf{f}}{\partial \mathbf{w}} (\hat{\mathbf{x}}_{k-1}^+, \mathbf{u}_{k-1}, 0) \quad (6.9)$$

$$H_k \equiv \frac{\partial \mathbf{h}}{\partial \mathbf{x}} (\hat{\mathbf{x}}_k^-, 0) \quad (6.10)$$

$$J_k \equiv \frac{\partial \mathbf{h}}{\partial \mathbf{v}} (\hat{\mathbf{x}}_k^-, 0) \quad (6.11)$$

The EKF linearizes around the previous state estimate  $\hat{\mathbf{x}}_{k-1}^+$  to generate an *a priori* estimate of the state vector  $\hat{\mathbf{x}}_k^-$  and covariance matrix  $P_k^-$ . This *a priori* covariance matrix  $P_k^-$  is used to determine the gain  $K_k$ . Applying this gain results in the best fit state estimate  $\hat{\mathbf{x}}_k^+$  for the current time step. The gain is also used to calculate the current best estimate of the *a posteriori* covariance matrix  $P_k^+$ .

In most attitude-determination applications, the states modeled by the EKF are the attitude parameters and angular rates. However, the attitude parameterization choice is not trivial. All attitude parameterizations of three dimensions contain a singularity in their kinematic differential equation, while three-dimensional attitude coordinates expressed in four or more dimensions have dependent parameters, which can result in a singular covariance matrix after an EKF update because numerical errors can cause one or more of the parameter constraints to be violated. The Multiplicative Extended Kalman Filter (MEKF) resolves this issue by using a combination of quaternions and a three-dimensional attitude parameterization; the former is used as a non-singular reference and the latter is used to tabulate the attitude error at each time step. The MEKF was

originally developed assuming the use of a rate gyro [54]. Alternatively, the rate gyro can be omitted by modeling the dynamics of the spacecraft. For the purposes of the MEKF, the PMAC attitude dynamics can be modeled by including only the bar magnet torque (which dominates all other external torques) and treating all other external torques as Gaussian process noise [11]. A disadvantage of this method is the inclusion of the mass moment of inertia matrix and the bar magnet strength in the dynamics model; these must be accurately measured before launch or fitted to on-orbit data to achieve satisfactory results.

We use the scaled Gibbs vector as the attitude error parameterization and follow the convention of Markley [54] to define the scaled Gibbs vector:

$$\mathbf{a}_g \equiv 2 \frac{\mathbf{q}}{q_0} \quad (6.12)$$

where  $\mathbf{q}$  is the vector part of the quaternion and  $q_0$  is the scalar part. The factor of two is included because it makes  $\mathbf{a}_g$  approximately equal to the yaw, pitch, and roll Euler angles for any rotation set, given that  $\mathbf{a}_g$  represents a small rotation. Thus, when the scaled Gibbs vector is used as the attitude error parameterization, the uncertainty estimate from the MEKF is directly applicable to satellite attitude uncertainty in yaw, pitch, and roll. Thus, the six-dimensional state vector for the PMAC MEKF is  $\mathbf{x} = [\mathbf{a}_g \ \boldsymbol{\omega}]^T$ . Although the reference quaternion is not technically a state, it is updated at the end of each filtering step as follows:

$$\hat{q}_k^+ = \delta q(\mathbf{a}_{g,k}) \otimes \hat{q}_k^- \quad (6.13)$$

where  $\delta q(\mathbf{a}_g)$  is the error quaternion given by

$$\delta q(\mathbf{a}_g) \approx \begin{bmatrix} \mathbf{a}_g/2 \\ 1 - \mathbf{a}_g^2/8 \end{bmatrix} \quad (6.14)$$

where the scalar-last quaternion convention is used. In practice, the *a posteriori* quaternion should be renormalized after each use of Equation 6.13 to combat numerical error buildup. Combining Equation 6.14 with the definition of a rotation matrix in terms of quaternions and assuming  $\mathbf{a}_g$  is a

small angle and ignoring higher-order terms yields the error rotation matrix. This rotation matrix definition is useful in determining the MEKF Jacobian matrices.

$$[R(\delta q(\mathbf{a}_g))] \approx [I_{3 \times 3}] - [\mathbf{a} \times] - \frac{1}{2}(a_g^2 [I_{3 \times 3}] - \mathbf{a}_g \mathbf{a}_g^T) \quad (6.15)$$

In order to determine the Jacobian matrices that define the behavior of the MEKF, the state dynamics must be determined. By combining Equation 6.13 with the kinematic differential equation for quaternions and the time-derivative of Equation 6.12, the time-derivative of the scaled error Gibbs vector can be obtained as shown below. The time-derivative of the angular velocity vector is given by Equation 2.1 as torqued by the bar magnet alone (Equation 2.6). The dynamics model used by the PMAC MEKF is thus:

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{\mathbf{a}}_g \\ \dot{\boldsymbol{\omega}} \end{bmatrix} = \begin{bmatrix} ([I_{3 \times 3}] + \frac{1}{4} \mathbf{a}_g \mathbf{a}_g^T)(\boldsymbol{\omega} - \hat{\boldsymbol{\omega}}) - \frac{1}{2}(\boldsymbol{\omega} + \hat{\boldsymbol{\omega}}) \times \mathbf{a}_g \\ [I]^{-1} (\mathbf{m} \times ([R(\delta q(\mathbf{a}_g))][R(q_{\text{ref}})]^T \mathbf{B} + \boldsymbol{\eta}_2) - \boldsymbol{\omega} \times [I] \boldsymbol{\omega} + \boldsymbol{\eta}_1) \end{bmatrix} \quad (6.16)$$

where  $\hat{\boldsymbol{\omega}}$  is the expected value of the angular velocity (as differentiated from the true angular velocity  $\boldsymbol{\omega}$ ); these are analogous to the difference between  $\hat{q}_k^-$  and  $\hat{q}_k^+$ , respectively. Note that we follow the method of Burton [11] and define the process noise vectors  $\boldsymbol{\eta}_1$  and  $\boldsymbol{\eta}_2$  as the inertial magnetic model error (including errors due to satellite position) and unmodeled external torques, respectively. With the state dynamics fully defined, the first two MEKF Jacobian matrices may be evaluated as follows:

$$F_k \equiv \frac{\partial \mathbf{f}}{\partial \mathbf{x}} (\hat{\mathbf{x}}_{k-1}^+, \mathbf{u}_{k-1}, 0) = \begin{bmatrix} -[\boldsymbol{\omega}_{k-1} \times] & [I_{3 \times 3}] \\ [I]^{-1} [\mathbf{m} \times][R(\hat{q}_{k-1}^+)]^T \mathbf{B}_{k-1} & [I]^{-1} (-[\boldsymbol{\omega}_{k-1} \times][I] + [I \boldsymbol{\omega}_{k-1} \times]) \end{bmatrix} \quad (6.17)$$

$$G_k \equiv \frac{\partial \mathbf{f}}{\partial \mathbf{w}} (\hat{\mathbf{x}}_{k-1}^+, \mathbf{u}_{k-1}, 0) = \begin{bmatrix} [0_{3 \times 3}] & [0_{3 \times 3}] \\ [I]^{-1} & [I]^{-1} [\mathbf{m} \times] \end{bmatrix} \quad (6.18)$$

Before evaluating Equations 6.10 and 6.11, first note that a body frame observation vector can be expressed using inertial data as  $\mathbf{h}(\mathbf{x}_k, \mathbf{v}_k) = [R(\delta q(\mathbf{a}_g))][R(\hat{q}_k^-)]^T \mathbf{b}_k + \mathbf{v}_k$  where  ${}^I \mathbf{b}_k$  is the observation

vector in the inertial frame at this step. With this full measurement vector model in hand, the final two Jacobians are evaluated as:

$$H_k \equiv \frac{\partial \mathbf{h}}{\partial \mathbf{x}} (\hat{\mathbf{x}}_k^-, 0) = \begin{bmatrix} [R(\hat{q}_k^-)^T \mathbf{b}_k \times] & [0_{3 \times n}] \end{bmatrix} \quad (6.19)$$

$$J_k \equiv \frac{\partial \mathbf{h}}{\partial \mathbf{v}} (\hat{\mathbf{x}}_k^-, 0) = [I_{6 \times n}] \quad (6.20)$$

where Equation 6.15 has been substituted prior to integration. The number of measurements at a given step determines the size of the Jacobians  $H_k$  and  $J_k$ , as well as the size of the measurement covariance matrix  $R$ . When using the filter, the observation vector is approximated assuming zero noise and using the expected value of the state to rotate an inertial, model-based observation at each step:

$$\mathbf{h}(\hat{\mathbf{x}}_k, 0) = [R(\hat{q}_k^-)^T] \mathbf{b}_k \quad (6.21)$$

To be clear, the *a priori* values  $\hat{q}_k^-$  and  $\hat{\omega}_k^-$  are obtained by numerically integrating the state dynamics given the *a posteriori* values from the previous step:

$$\begin{bmatrix} \hat{q}_k^- \\ \hat{\omega}_k^- \end{bmatrix} = \begin{bmatrix} \hat{q}_{k-1}^+ \\ \hat{\omega}_{k-1}^+ \end{bmatrix} + \int_{t_{k-1}}^{t_k} \begin{bmatrix} \dot{q} \\ \dot{\omega} \end{bmatrix} dt. \quad (6.22)$$

Note that the scaled error Gibbs vector  $\mathbf{a}_g$  is not iterated;  $\mathbf{a}_g$  is defined as the error between the *a priori* and *a posteriori* estimates. The expected value of the scaled error Gibbs vector expected value is always zero, but it does have a non-zero *a posteriori* estimate set by Equation 6.6. This is why the quaternion *a posteriori* update (Equation 6.13) is always the final action of each filtering time step.

The PMAC MEKF derivation assumes that the error Gibbs vector represents a small angle rotation (usually defined as  $<5^\circ$ ). As a result, trust in the MEKF-output states and uncertainties should be tentative if components of the error Gibbs vector are consistently  $>5^\circ$ . The PMAC

MEKF filter is now fully defined. However, the filter must be tuned before its outputs may be trusted.

## 6.2 Filter Tuning

Although the measurement noise  $R$  can be usually be determined via sensor testing, the process noise  $Q$  is more difficult to set. Tuning is the process of varying  $Q$  and  $R$  as needed until the filter performance is acceptable. After tuning, the uncertainty given by the filter covariance should match the error of the state estimates. There are two tuning methods: simulating noisy input to the filter and empirical tuning using residuals; the CSSWE mission used both methods.

### 6.2.1 Simulation-based Filter Tuning

The simulation uses the SGP4 orbital propagator [79], the CSSWE TLE, an initial attitude, satellite properties, and environmental torques models to numerically determine the “true” satellite motion. The simulation shares many similarities with the one developed in Chapter 8. External torques which the MEKF does not model (gravity-gradient, drag, and solar pressure) are used in this “truth” simulation to determine if estimating these disturbance torques as Gaussian noise is acceptable. Gaussian noise of 15nT is added to the IGRF model before calculating the magnetic torque at each time step; this value models good orbital position knowledge with a quiet sun. Noise is also added to simulated body-frame measurements; the  $1\sigma$  standard deviation of the photodiodes is set to  $1.6\mu\text{A}$  which reflects the  $50\text{ W}\cdot\text{m}^{-2}$  uncertainty in the albedo model used to correct the photodiode output. The  $1\sigma$  magnetometer standard deviation is set to 400nT as it is based on the encountered  $\mathbf{B}$ -field magnitude error (see Figure 4.5). Noisy body-frame measurements generated by the simulation are then filtered by the MEKF. This filtering is repeated with a sensible range of  $Q$  (based on expected levels of the external torques which are not modeled and the input IGRF noise) until the uncertainty bounds correctly represent the errors (difference between “truth” and the MEKF-output) in filtered attitude and angular velocity.

Figures 6.1 and 6.2 show the properly-tuned MEKF output given simulated measurements

Figure 6.1: Angular error as output by the MEKF after tuning using a simulation of the attitude dynamics. The error is the difference between the simulated attitude and the MEKF fitted attitude generated from measurements corrupted with Gaussian noise. The red line shows the  $3\sigma$  angular uncertainty, which should bound 99.73% of the angular error.

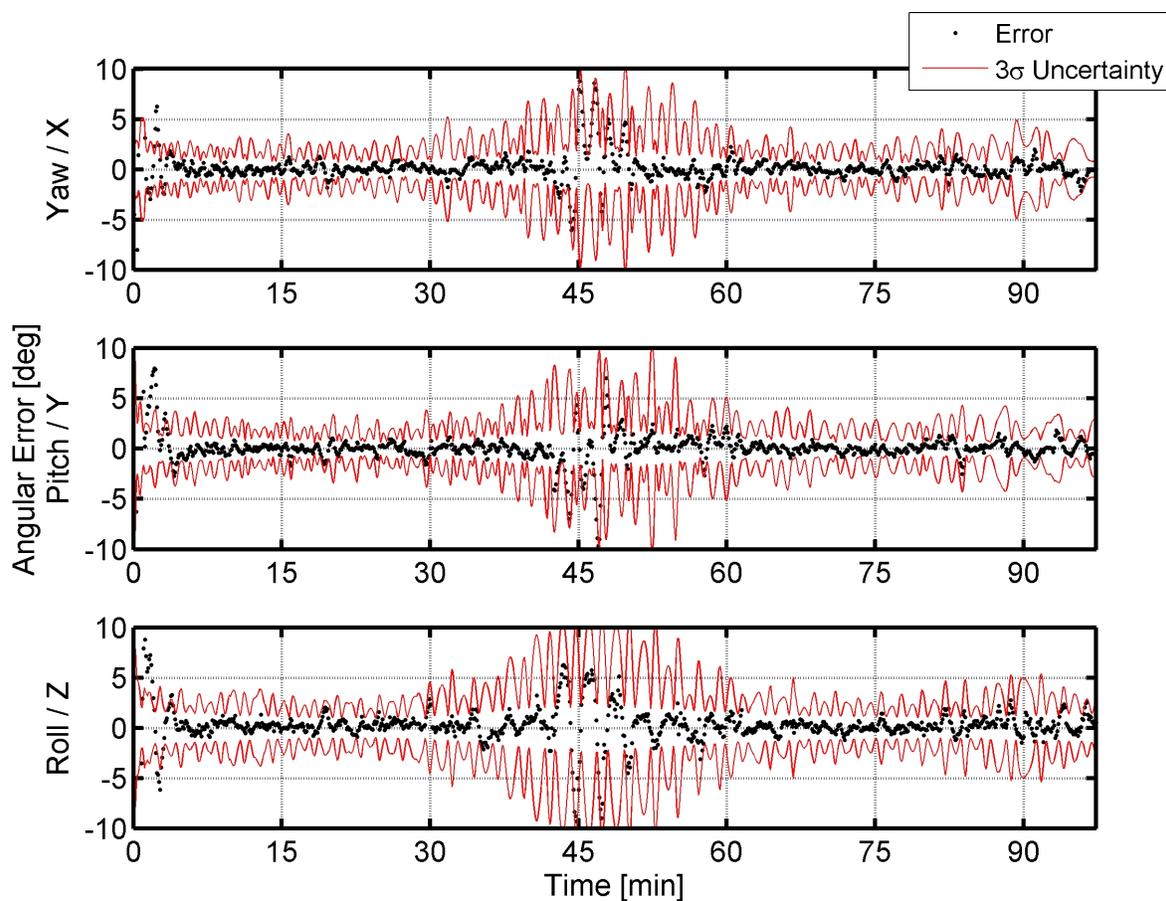
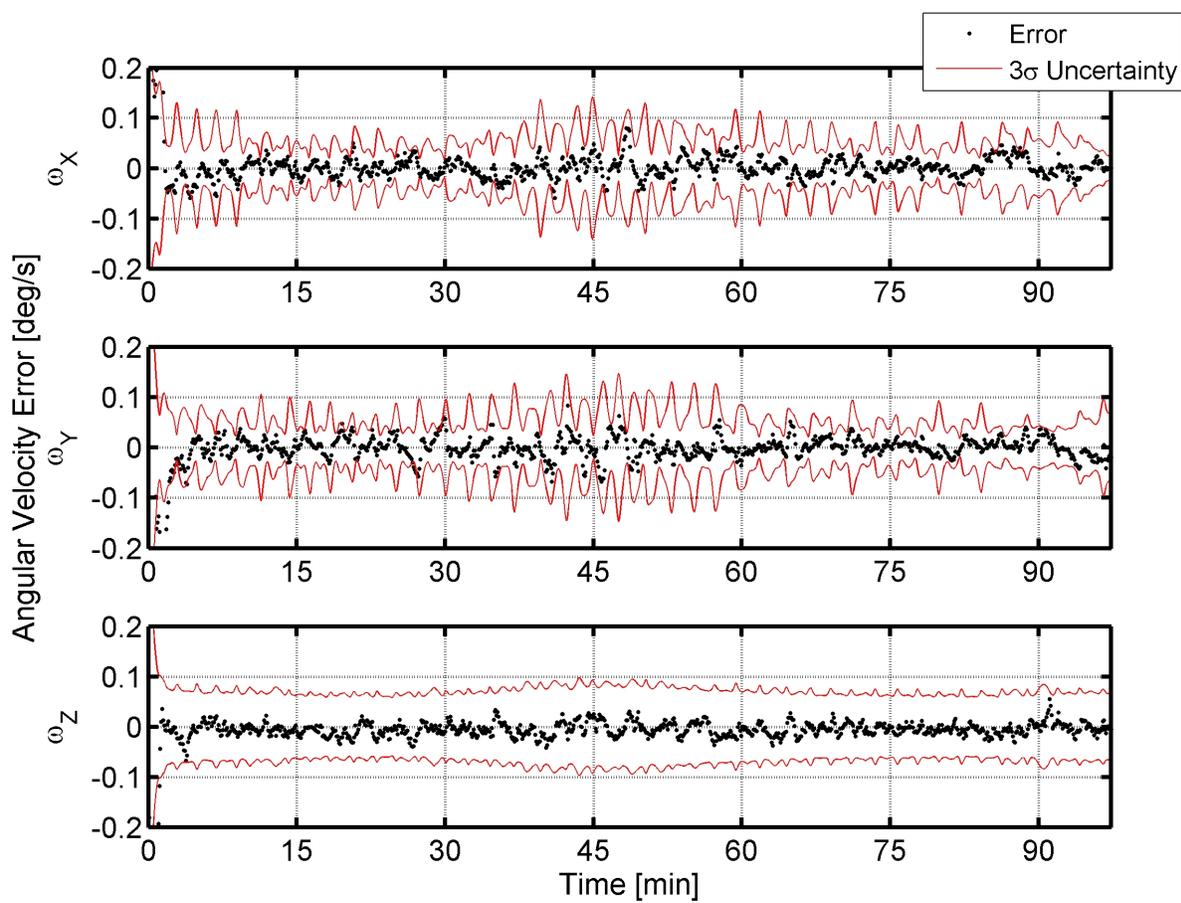


Figure 6.2: Error in angular velocity as output by the MEKF after tuning using a simulation of the attitude dynamics. The error is the difference between the simulated angular velocity and the MEKF fitted angular velocity generated from measurements corrupted with Gaussian noise. The red line shows the  $3\sigma$  angular velocity uncertainty, which should bound 99.73% of the angular velocity error.



over a typical early-mission CSSWE orbit. Although each simulation is slightly different due to its stochastic nature, the  $3\sigma$  bounds reliably encase roughly 99% of the angular error. The short-term spikes in uncertainty are when the satellite is aligned with the local magnetic field; this alignment reduces the signal to noise ratio on the other axes of the magnetometer, resulting in a less certain attitude estimate. These uncertainty spikes are especially pronounced from 30 to 60 minutes, when the photodiode measurements are ignored during eclipse.

As described in Section 4.3.2, CSSWE uses a magnetometer and four photodiodes to observe its attitude. Figures 6.3 and 6.4 show the measurement residuals for each magnetometer axis and photodiode. The simulation shows that 99% of the measurement residuals are also within the  $3\sigma$  uncertainty bounds, which are calculated using both the assumed standard deviations and the MEKF state uncertainty output.

Figure 6.5 shows the components of the scaled error Gibbs vector for each step in the simulated. As described above, the MEKF output may be trusted if these errors remain below  $5^\circ$ . After a settling time of about seven minutes, the error angles remain well below the  $5^\circ$  mark. However, the error angles do approach  $5^\circ$  during the eclipse period (30 to 60 minutes). Thus, the filter may be operating close to its performance limits during eclipse. On-orbit data was used for the next stage in filter tuning.

### 6.2.2 Empirical Filter Tuning

When working with the on-orbit CSSWE attitude data, a timespan of interest was set from the first on-orbit measurement (September 14, 2012, at 00:51:50 UTC) to midnight the day before a latch-up resulted in degraded photodiode measurements (October 14, 2012 at 00:00:00 UTC). This month of time has much variation, from initial tumbling to settling, from safe mode to science mode, and from relatively quiet geomagnetic conditions to a significant storm.

The empirical tuning was complicated further by uncertainty in the physical properties of the satellite. As explained in Section 6.1, the bar magnet moment  $\mathbf{m}$  and the inertia matrix  $[I]$  are key factors in the performance of the PMAC MEKF. The inertia matrix  $[I]$  was numerically calculated

Figure 6.3: Magnetometer measurement residuals from the simulated-input PMAC MEKF output. The residual is the difference between the measured value and the inertial model rotated into the body frame. The red line shows the  $3\sigma$  angular uncertainty, which should bound 99.73% of the residuals.

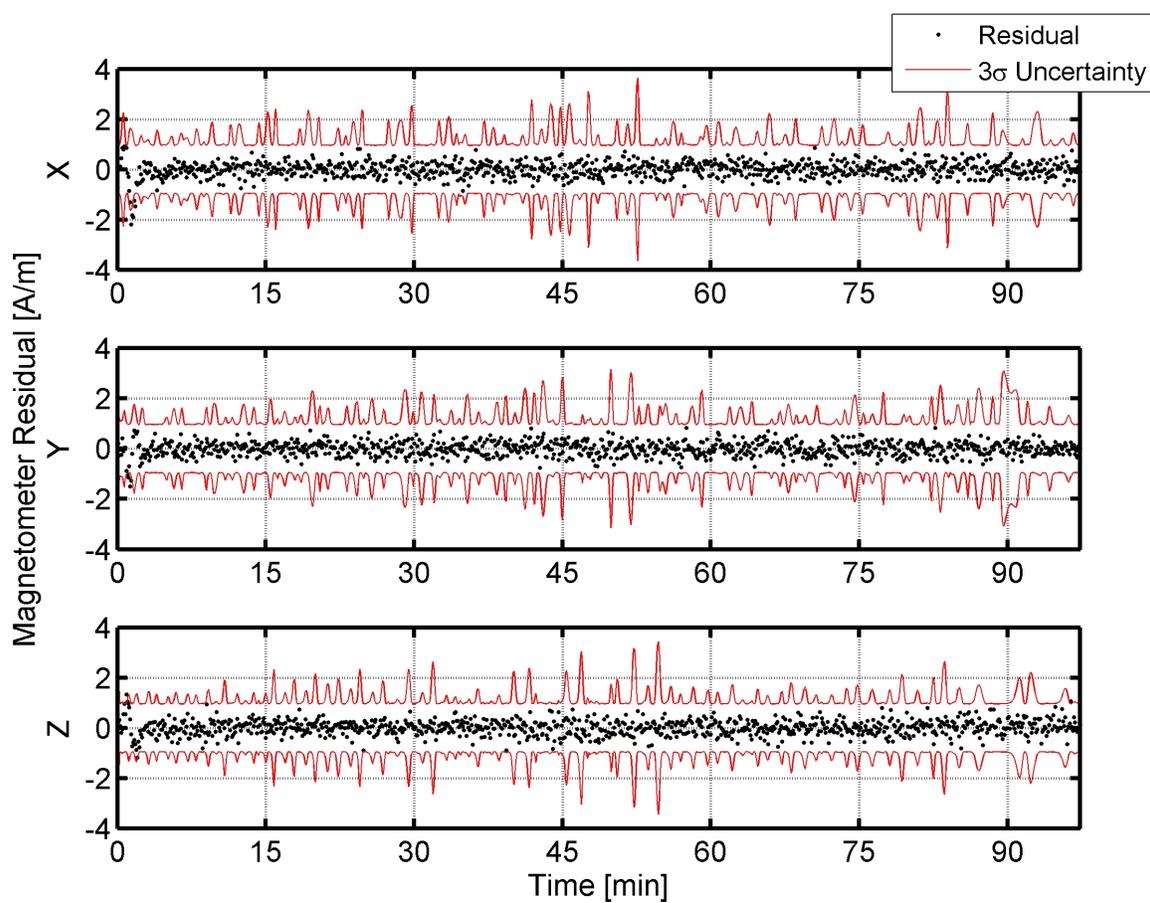


Figure 6.4: Photodiode measurement residuals from the simulated-input PMAC MEKF output. The residual is the difference between the measured value and the inertial model rotated into the body frame. The red line shows the  $3\sigma$  angular uncertainty, which should bound 99.73% of the residuals.

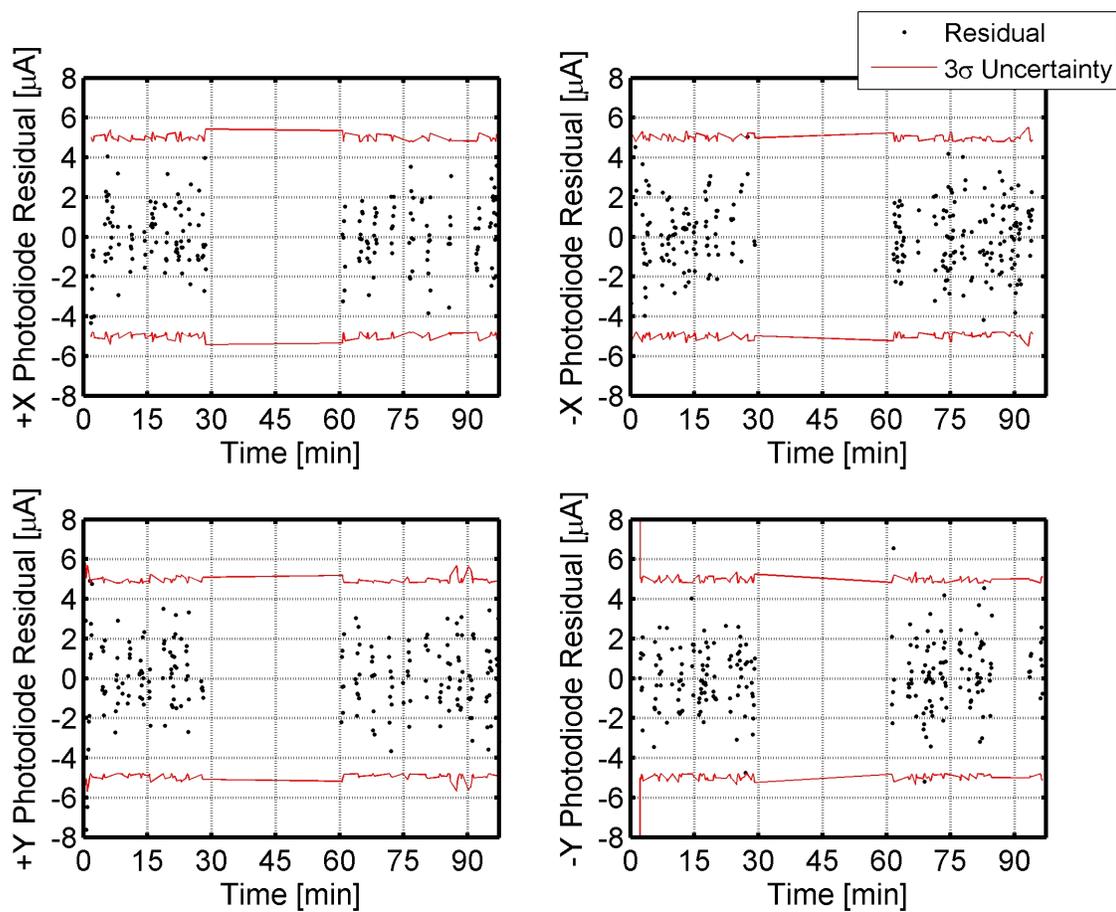
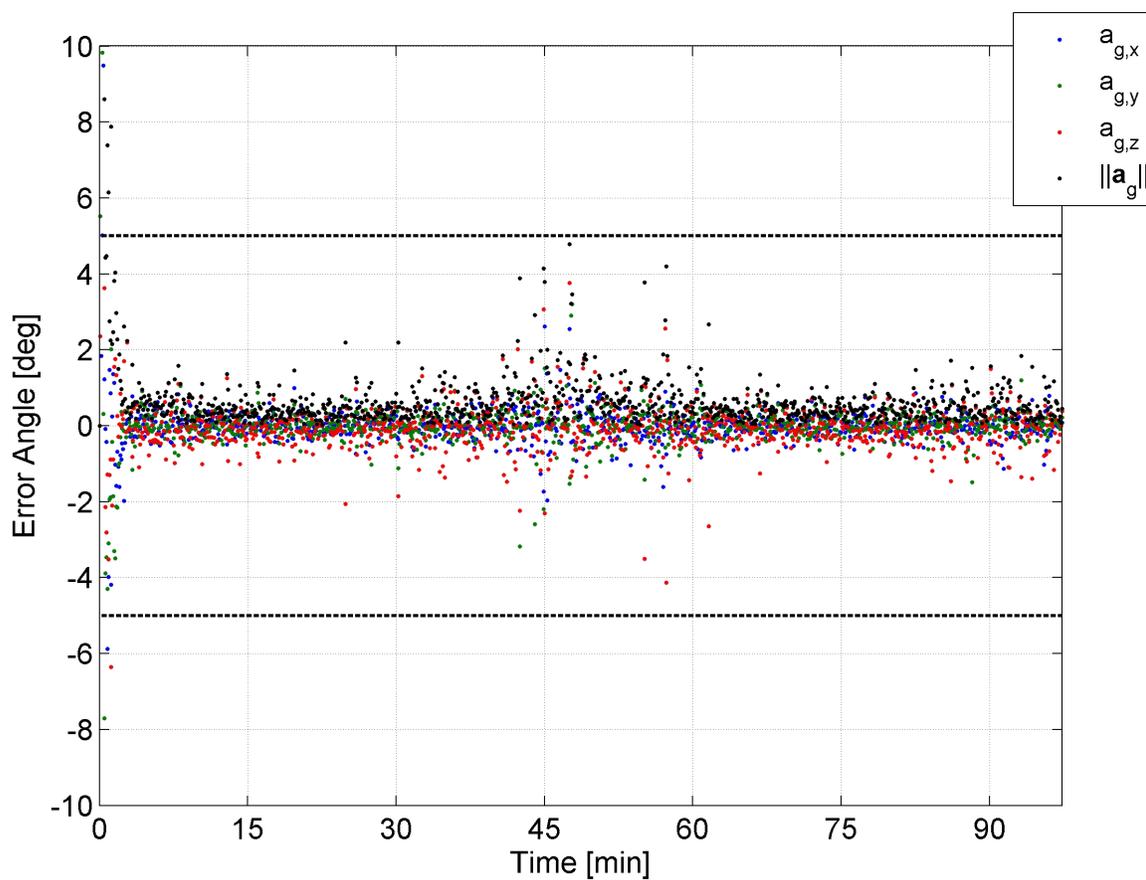


Figure 6.5: The scaled error Gibbs vector values from the simulated-input PMAC MEKF. When the error angle is  $<5^\circ$ , the MEKF output may be trusted.



from the spacecraft solid model. The uncertainty of this estimate is based on the completeness of the solid model; in the case of CSSWE, the  $[I]$  uncertainty is estimated to be roughly 5%. The magnetic moment of the bar magnet was measured using the technique described in Section 7.2 which shows an uncertainty of 2%. The MEKF simulation (described in the previous section) was used to test filter performance given incorrect values of  $[I]$  and  $\mathbf{m}$  at the maximum expected uncertainty; the results were not satisfactory.

However, both  $[I]$  and  $\mathbf{m}$  should not change after all deployables are in their operational configuration; this makes on-orbit calibration ideal for these values. Constrained nonlinear optimization was used to determine the  $\mathbf{m}$  which minimized the filter measurement residuals. The fit was performed using early mission CSSWE data, while the spacecraft was still covering the majority of the attitude sphere. The inertia matrix was not fitted because it was found that  $\mathbf{m}$  and  $[I]$  are not separable from the viewpoint of the CSSWE attitude response. Thus, the solid model  $[I]$  and the fitted  $\mathbf{m}$  are used.

However, because the fit  $\mathbf{m}$  is based on minimizing the filter measurement residuals, the empirical tuning occurred simultaneously to the fit. A five hour timespan in the early mission is used to simultaneously fit  $\mathbf{m}$ , the photodiode and magnetometer alignment angles, and the photodiode scale factor. Longer fit times were attempted, but the results were degraded by the photodiode scale factor changes over their duration. The MEKF assumes photodiode  $1\sigma$  standard deviations of  $1.6 \mu\text{A}$ , as outlined in Section 6.2.1. Each magnetometer axis  $1\sigma$  standard deviation is set equal to the  $\|\mathbf{B}\|$  error (measured vs. IGRF); this value varies over the timespan of interest as shown in Figure 4.5. The magnetometer X-axis contains a systematic error which is believed to be due to its proximity to a current-carrying wire within the satellite. Unfortunately, the time-varying calibration using the 10 minute average system current housekeeping telemetry lacks the temporal resolution necessary to adequately remove this error. This error is thus reflected in the tuned  $R$  matrix.

Starting with the simulation-based value of  $Q$  and expected level of  $R$ , the empirical filter-based fit is performed multiple times with various  $Q$  and  $R$  until the approximately 99% of the

measurement residuals are within the  $3\sigma$  measurement uncertainty bounds. Empirical tuning shows that the best performance of the filter occurs when the magnetometer X-axis element of  $R$  is set to twice the nominal value (measured vs. IGRF  $\|\mathbf{B}\|$ ) throughout the timespan of interest. Figure 6.6 shows the magnetometer position and orientation on the REPTile board within the satellite. The X-axis of the magnetometer is parallel to the currents that are generated by the REPTile ground plane.

Figures 6.7 and 6.8 show the tuned magnetometer and photodiode measurement residuals, respectively, for the approximately five hour early mission empirical fit timespan. As explained above, the magnetometer X-axis shows a systematic error which has been accounted for by doubling the standard deviation of that sensor. The photodiode residuals behave as expected. In both sensors, approximately 99% of the measurement residuals reside within the  $3\sigma$  error bounds, proving that the empirical data has been properly fit. However, the decreased confidence in the magnetometer X-axis exacerbates the problem of high error angles during eclipse times, as shown in Figure 6.9. Although the figure shows about a dozen data points above the  $5^\circ$  threshold, the filter remains below this cutoff for the vast majority of the time. We have confidence in the filtered data during the timespan of interest because of the results of both tuning procedures.

### 6.3 CSSWE Attitude Determination

The tuned MEKF was applied to the first month of CSSWE on-orbit data, before the photodiodes were compromised by an on-orbit anomaly, as described in Section 4.4. The filter output is shown over this timespan, followed by a validation of the results.

#### 6.3.1 Filter Output

The output from the MEKF using the on-orbit data is split into multiple sections. The first section shows data over the entire timespan of interest. The next two sections show 100 minute datasets representative of the PMAC performance before and after settling occurs.

Figure 6.6: A picture of the interior of the CSSWE as captured during final integration. The magnetometer is located on the corner of the REPTile electronics board (bottom left of the image); the system axes are indicated. The X-axis is parallel to the currents generated by the REPTile ground plane.

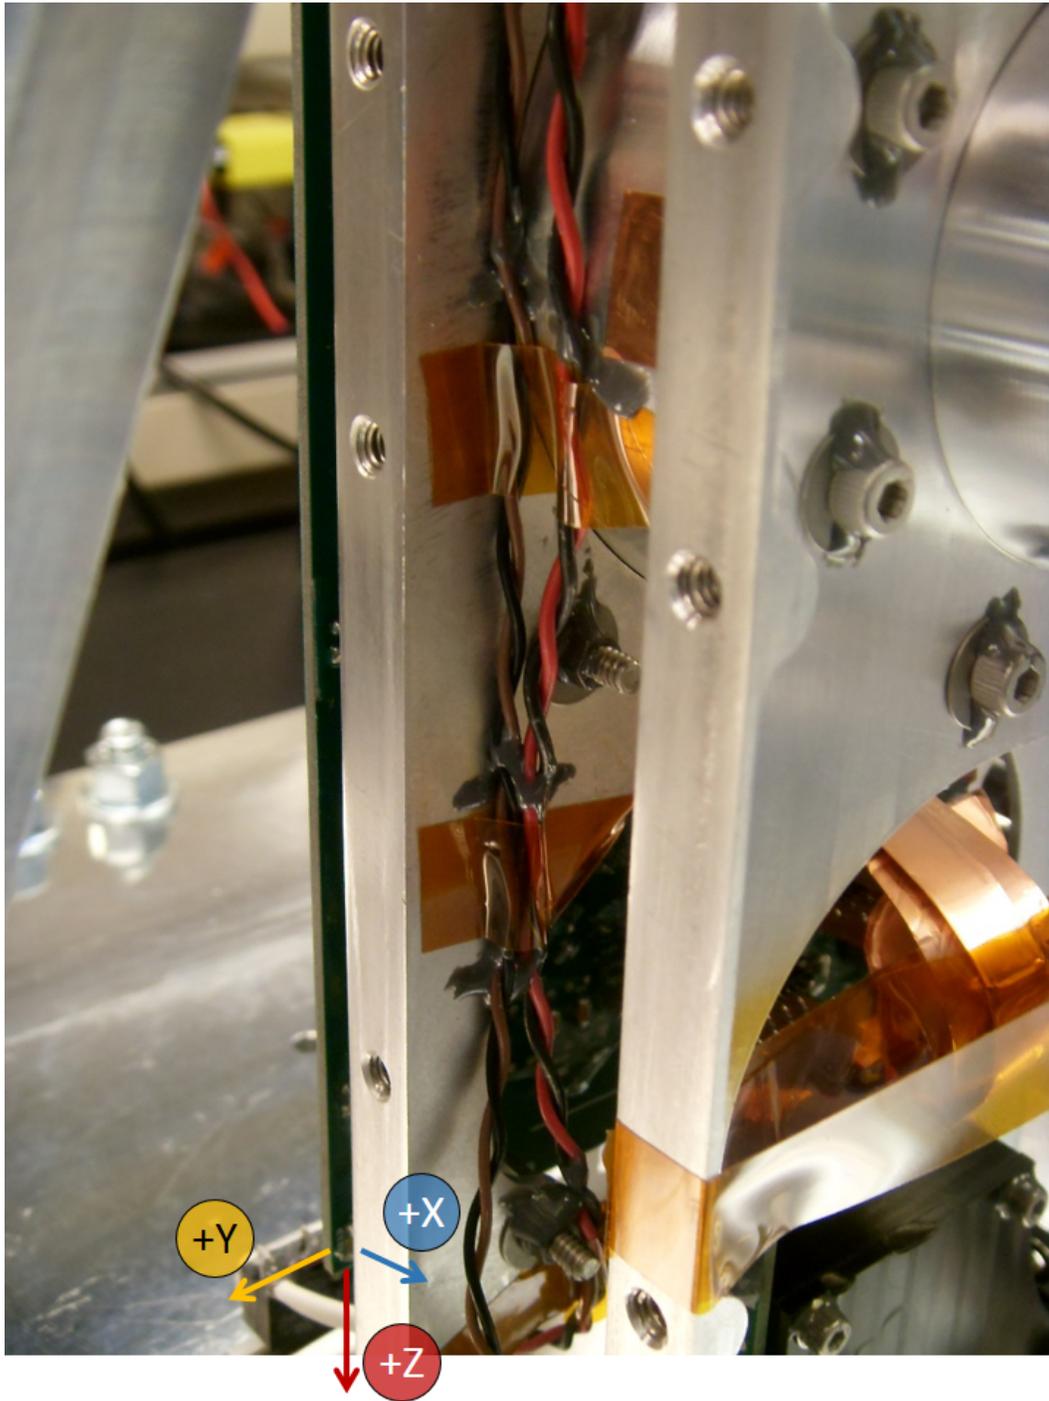


Figure 6.7: Magnetometer measurement residuals from the early mission on-orbit PMAC MEKF output. The residual is the difference between the measured value and the inertial model rotated into the body frame. The red line shows the  $3\sigma$  angular uncertainty, which should bound 99.73% of the residuals. The systematic error in the X-axis magnetometer output is believed to be due to system currents which cannot be removed using the available telemetry. The expected standard deviation of the X-axis magnetometer data (within the  $R$  matrix) is increased to twice the nominal level to account for this error.

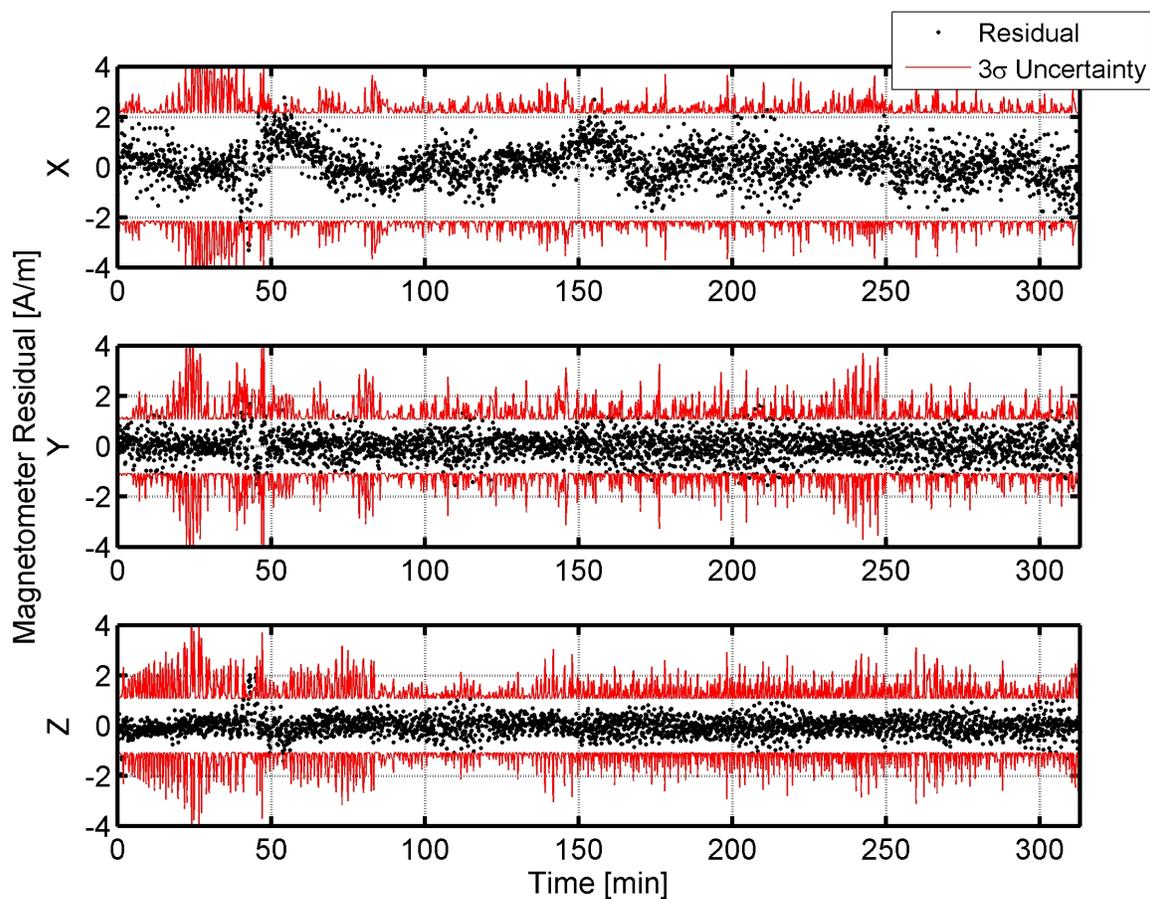


Figure 6.8: Photodiode measurement residuals from the early mission on-orbit PMAC MEKF output. The residual is the difference between the measured value and the inertial model rotated into the body frame. The red line shows the  $3\sigma$  angular uncertainty, which should bound 99.73% of the residuals.

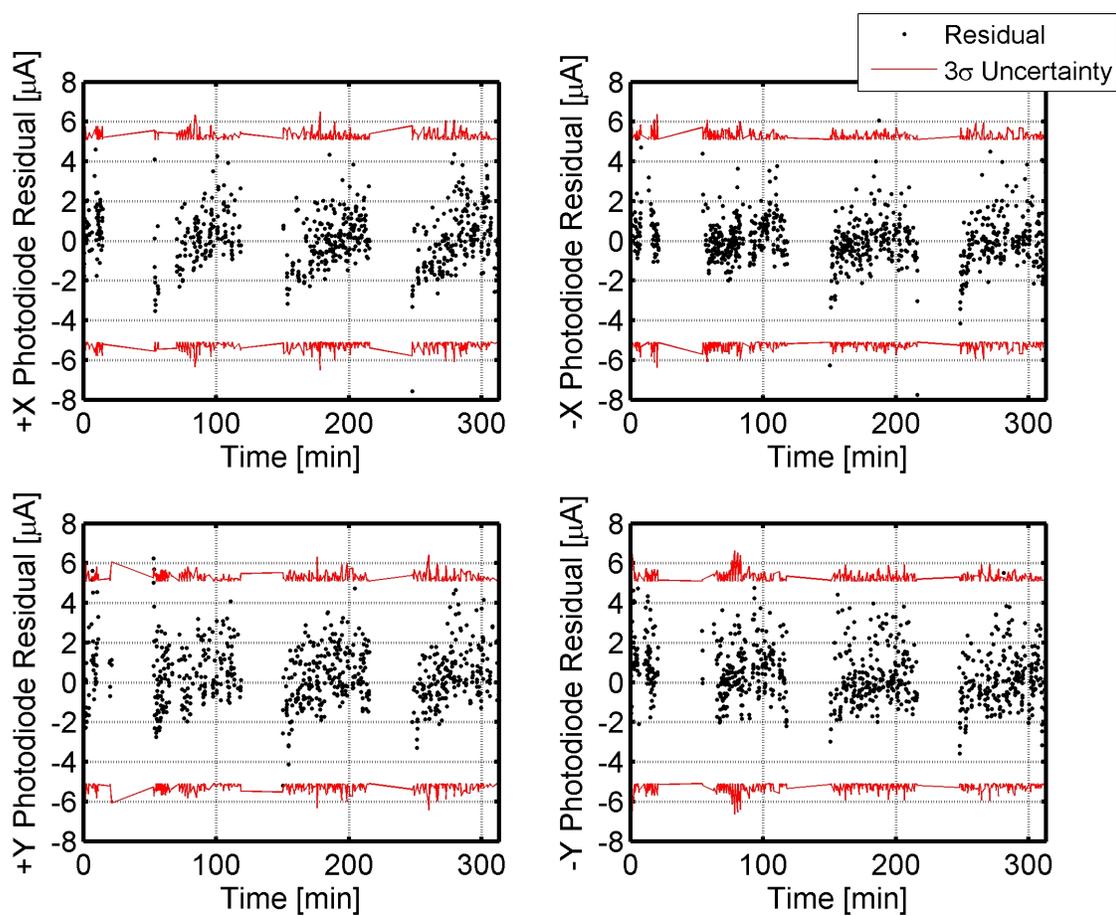
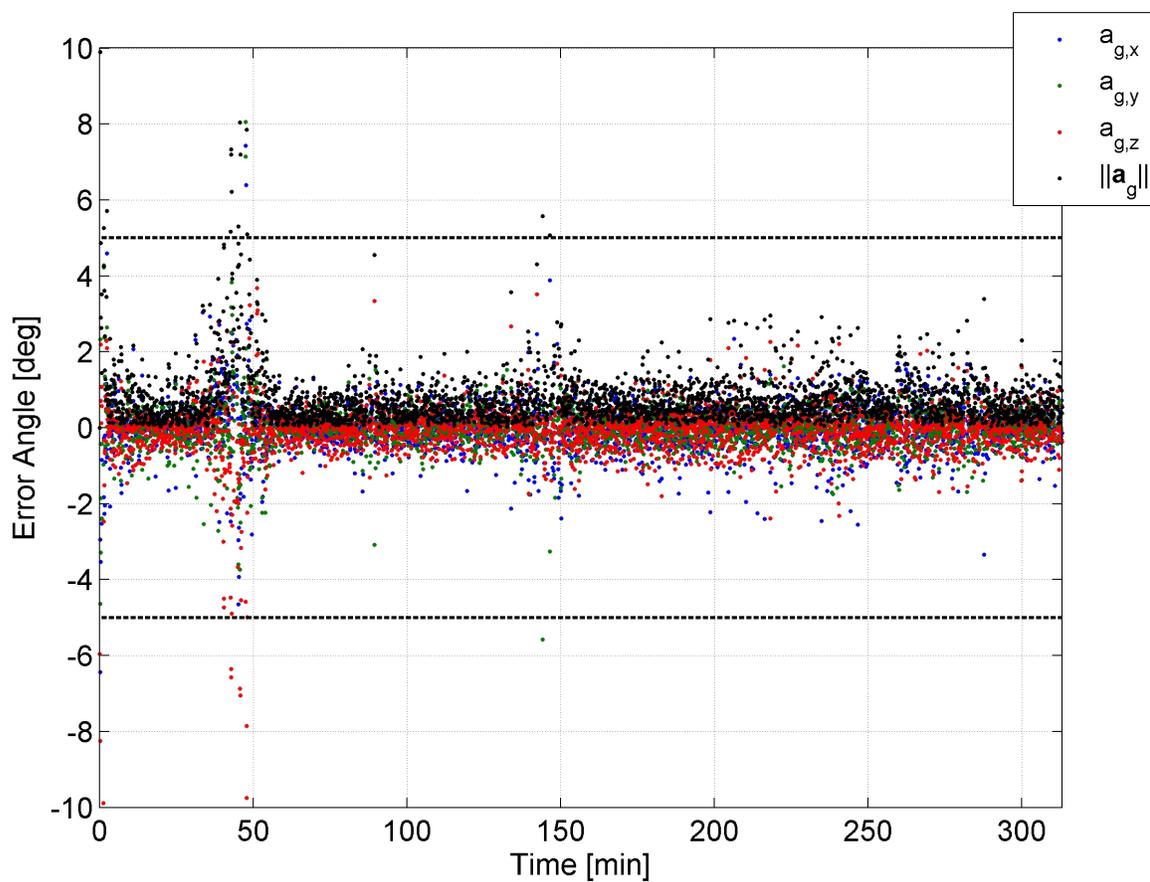


Figure 6.9: The scaled error Gibbs vector values from the early mission on-orbit PMAC MEKF. When the error angle is  $<5^\circ$ , the MEKF output may be trusted.



### 6.3.1.1 First 30 Days On-Orbit

Figure 6.10 shows the angular velocity components throughout the timespan of interest. The  $3\sigma$  uncertainty bounds are included in the plot in red but there is no visible difference from the fitted values at the zoom level shown. The spikes in the fitted angular velocity and uncertainty bounds are due to gaps in the housekeeping data which is used to correct the magnetometer readings or gaps in the attitude data itself. The  ${}^{\mathcal{B}}X$  and  ${}^{\mathcal{B}}Y$  components of the angular velocity dampen within the first week, and remain at low values throughout the month. The roll about the  ${}^{\mathcal{B}}Z$  axis varies greatly within the first few days, then settles somewhat. The roll continues to evolve within a few degrees per second with a zero crossing about 13 days after launch. The zero roll rate causes the satellite exterior to experience a wider range of temperatures as detailed in Section 6.3.2.2.

Figure 6.11 shows the  $\beta$  angle (between  ${}^{\mathcal{B}}Z$  and the  $\mathbf{B}$ -field) with the  $3\sigma$  uncertainty bounds. CSSWE settled to within  $15^\circ$  of the  $\mathbf{B}$ -field within seven days of launch. The increased magnitude of the  $3\sigma$  uncertainty bound spikes in the middle of the dataset are likely related to a decreased roll rate combined with increased eclipse periods. As shown in Figure 4.6, the number of photodiodes visible to the sun (and thus the number of measurements available to the MEKF) is dependent on the roll angle. When the roll rate is slowed, the satellite has a longer timespan in which only one photodiode is viable for use. As the satellite eclipse period increases (see Figure 6.15), the photodiodes are not used for a longer period and the uncertainty bounds grow to a higher maximum as the MEKF is using less observations for a longer time.

The attitude estimate is further degraded due to the satellite orientation with respect to the local magnetic field. After the attitude has settled, the two magnetometer axes perpendicular to the  $\mathbf{B}$ -field have a much lower signal-to-noise ratio and the MEKF performance degrades as a result. This could be avoided by installing the magnetometer significantly off-axis from the satellite body frame.

The system energy evolution is a good “sanity check” for any dynamics problem. The rotational energy of a PMAC satellite may be split into the kinetic and potential energy. The rotational

Figure 6.10: The body to inertial angular velocity vector components are shown over the first month of CSSWE on-orbit operations. The  $3\sigma$  uncertainty bounds are shown in red behind the fit data; their inclusion makes a negligible difference at this zoom level.

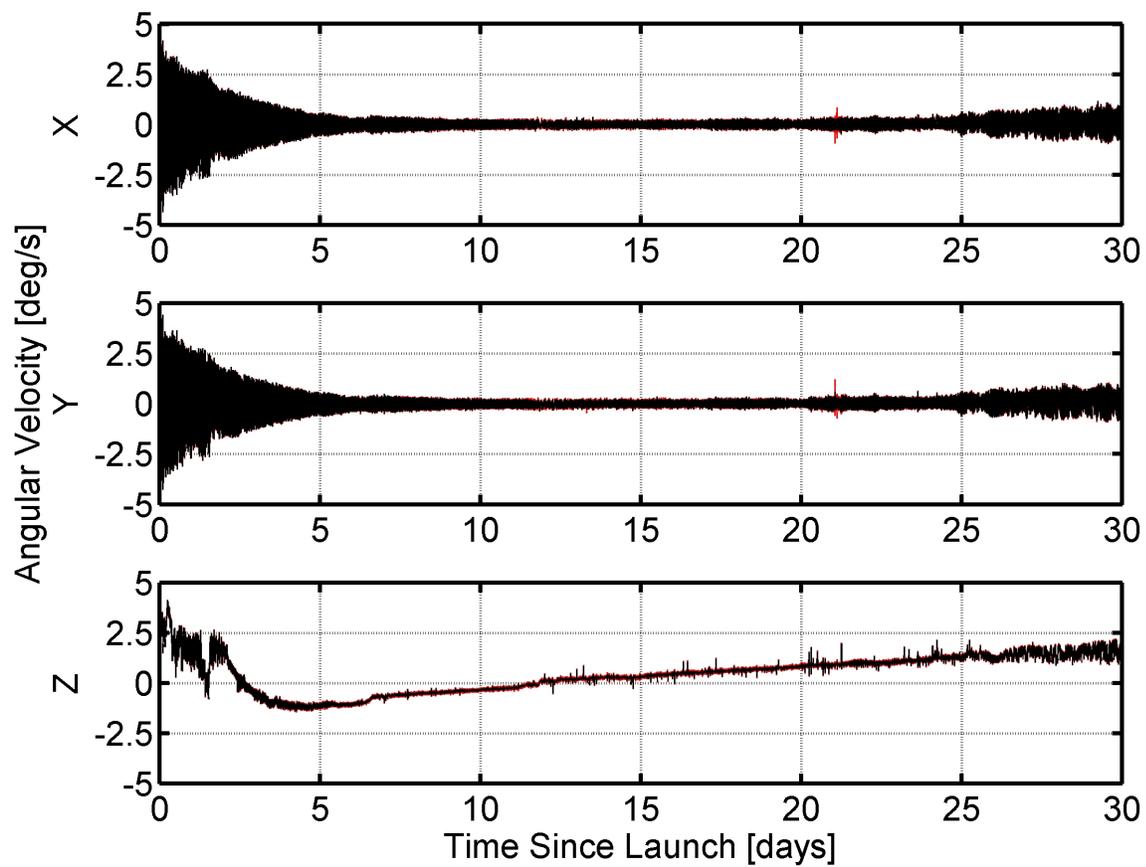


Figure 6.11: The angle  $\beta$  between the local magnetic field and the CSSWE  $+^{\beta}Z$  axis is shown for the first month on-orbit with  $3\sigma$  uncertainty bounds included in red. The  $15^{\circ}$  beta angle threshold indicating post-settling is denoted with a solid green line. The increased beta angles late in the dataset may be erroneous output due to satellite position error as explained in Section 4.3.2.3. The increased uncertainties in the middle of the dataset are likely due to decreased roll rates combined with increased eclipse times.

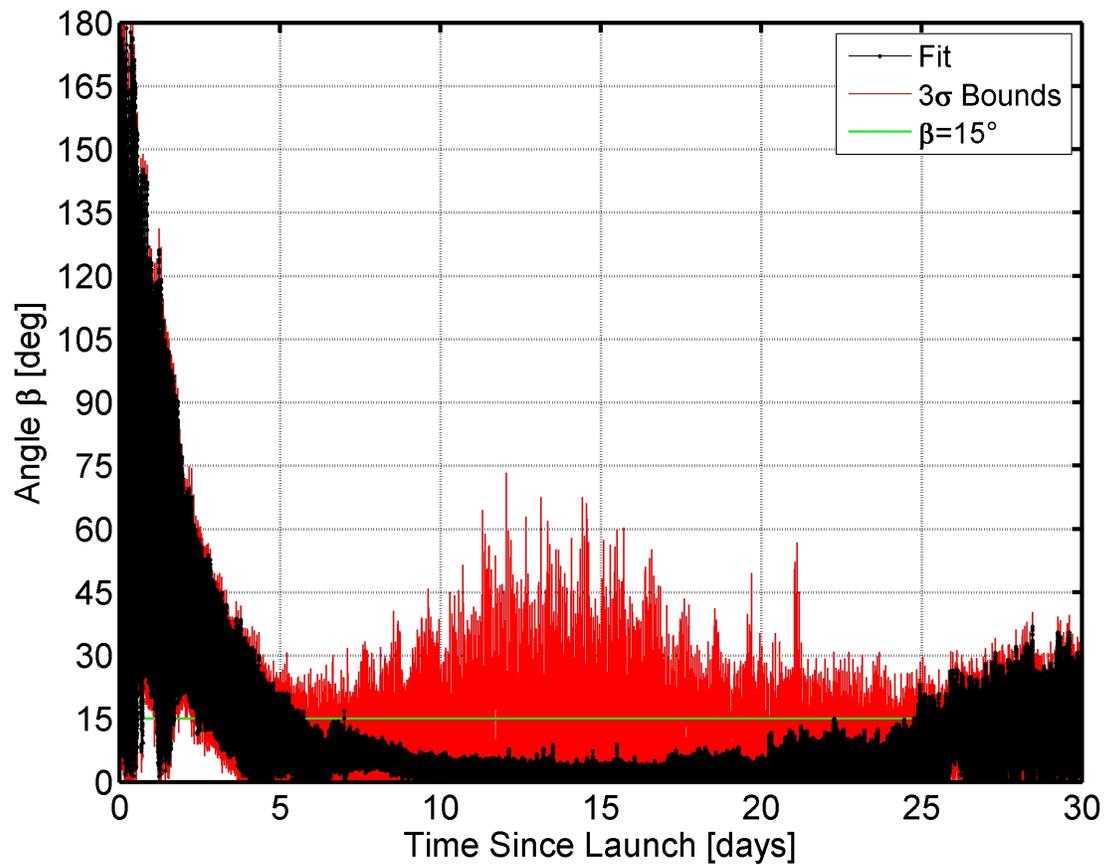


Figure 6.12: The rotational energy of the satellite over the first month on-orbit. The total energy is shown in black at bottom, while the top shows the kinetic (red) and magnetic potential (blue) energy components.

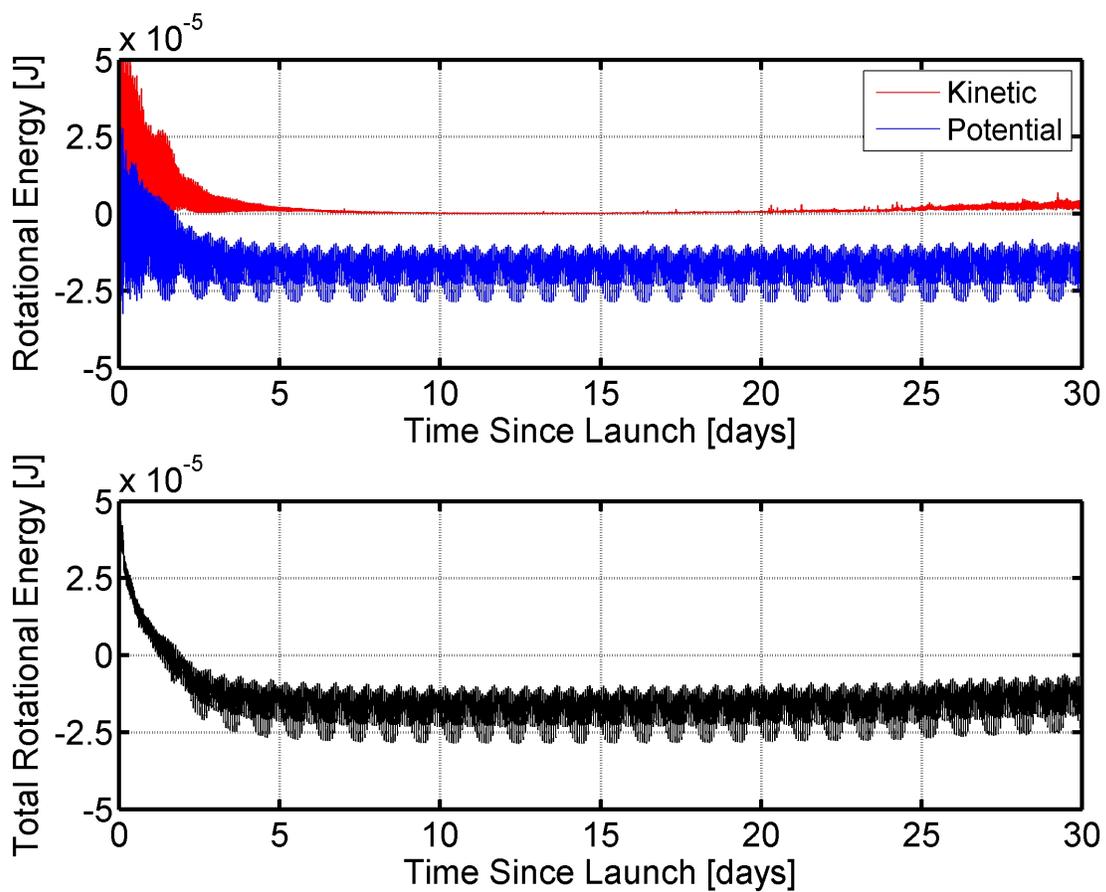
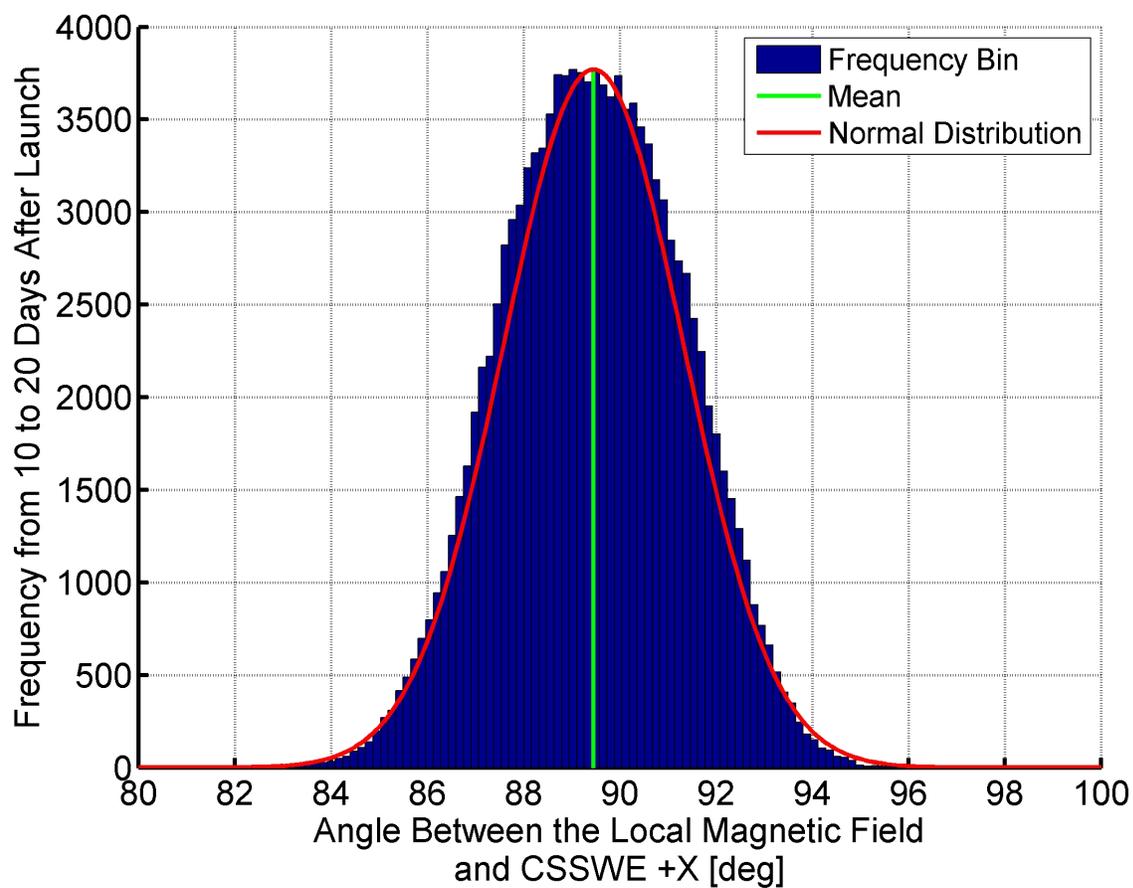


Figure 6.13: The histogram of the angle between the local magnetic field and the CSSWE + $\beta$ X (REPTile aperture) axis. The histogram covers a ten day period starting ten days after launch (after the  $\beta$  angle has settled as shown in Figure 6.11). The mean angle of  $89.5^\circ$  is shown in green and a fitted normal distribution with a standard deviation of  $1.87^\circ$  is shown in red.



kinetic and potential energy equations are repeated below.

$$T_K = \frac{1}{2} \boldsymbol{\omega}^T [I] \boldsymbol{\omega} \quad (5.2)$$

$$T_P = -\mathbf{m} \cdot \mathbf{B} \quad (5.3)$$

The rotational potential energy is due to magnetic material within the satellite. Note that Equation 5.3 defines zero potential energy as when the bar magnet is perpendicular to the  $\mathbf{B}$ -field. The satellite possesses negative potential energy when the magnetic moment vector is less than  $90^\circ$  away from the  $\mathbf{B}$ -field. Note that the potential energy is directly related to the  $\mathbf{B}$ -field magnitude, which varies over the CSSWE orbit. Figure 6.12 shows the kinetic, potential, and total rotational energy over the first month on-orbit. The hysteresis rods decrease the rotational kinetic energy to nearly zero within the first week after launch. The potential energy decreases as the satellite aligns but the settled behavior echoes the orbital variations in the  $\mathbf{B}$ -field magnitude. The increased jitter in the kinetic energy after day twenty is believed to be an effect of the increased error in satellite position estimation (Figure 4.8) which similarly affects the beta angle and angular velocity estimates.

After settling, the angle between the REPTile aperture ( $+^B X$ ) and the  $\mathbf{B}$ -field should be close to  $90^\circ$ . Figure 6.13 is a histogram of this REPTile alignment angle from days 10 to 20 after launch, after the attitude has settled and before the filter performance has degraded. As shown, the REPTile alignment angle relative to the magnetic field follows a normal distribution with a mean of  $89.5^\circ$  and a standard deviation of  $1.87^\circ$ . As mentioned in Chapter 4, this is advantageous as it provides a higher signal-to-noise ratio for the REPTile instrument.

Figure 6.14 shows MEKF scaled error Gibbs vector magnitude statistics over the first month on-orbit. The daily mean, standard deviation, and percentage of magnitudes greater than  $5^\circ$  are shown; each is a relative measure of the MEKF performance. By each measure, the MEKF performs best in the early mission once the initial angular velocity has settled somewhat. As noted

in Section 4.3.2, both the magnetometer and photodiodes require on-orbit calibration to achieve acceptable performance. The majority of the calibration parameters are set using early mission data; this is why the MEKF performance degrades after the first few days on-orbit. The REPTile instrument was activated for the first time on day twenty; this event sharply increases the error angle magnitude because the magnetometer is located on the REPTile electronics board (see Figure 4.4). From day 20 to 25, the REPTile detectors are being activated one by one. REPTile completes its commissioning on day 25; after this point, the magnetometer calibration better compensates for REPTile currents.

Figure 6.15 shows the daily maximum, mean, and minimum temperatures of the Command and Data Handling (C&DH) board in the interior of CSSWE. The orbit percent illuminated by the sun is also shown. It is easy to see that the interior temperature reflects the satellite insolation time. As explained in Section 4.3.2.1, the C&DH temperature is used as a proxy for the magnetometer temperature, which is responsible for the most significant time-varying magnetometer scaling and offset errors. The temperature correction is essentially extrapolating based on the first three days on-orbit; the magnetometer error can be expected to increase as the temperatures differ from the first three days. The magnetometer temperature has a nonlinear relationship to the magnetometer error which has consequences for the MEKF performance (Figure 6.14). Figure 6.15 shows a second source of MEKF uncertainty: the orbit eclipse period reaches its maximum near day 13, when the daily percent error angle magnitude is near its maximum; this is the effect of losing the photodiode measurements for a longer period.

### 6.3.1.2 Early Mission On-orbit Performance

Figure 6.16 and 6.17 show the  $\beta$  angle and angular velocity components with their associated  $3\sigma$  bounds for a 100 minute period four hours after orbit insertion. Both plots show a satellite that has not yet settled to pointing parallel to the  $\mathbf{B}$ -field. The  $\beta$  angle is still varying up to almost  $180^\circ$ . The roll rate is about  $4^\circ/\text{s}$  and the pitch and yaw rates are exchanging at  $\pm 3^\circ/\text{s}$  due to gyroscopic torques induced by the satellite inertia matrix asymmetry. The fit uncertainty does not extensively

Figure 6.14: The daily mean and standard deviation of the scaled error Gibbs vector magnitude  $\|\mathbf{a}_g\|$  are shown in blue squares and green circles, respectively. The daily percent of error angle magnitudes  $\|\mathbf{a}_g\|$  which are greater than  $5^\circ$  is shown using the red triangles. Each dataset is a relative measure of the MEKF performance over the first thirty days on-orbit.

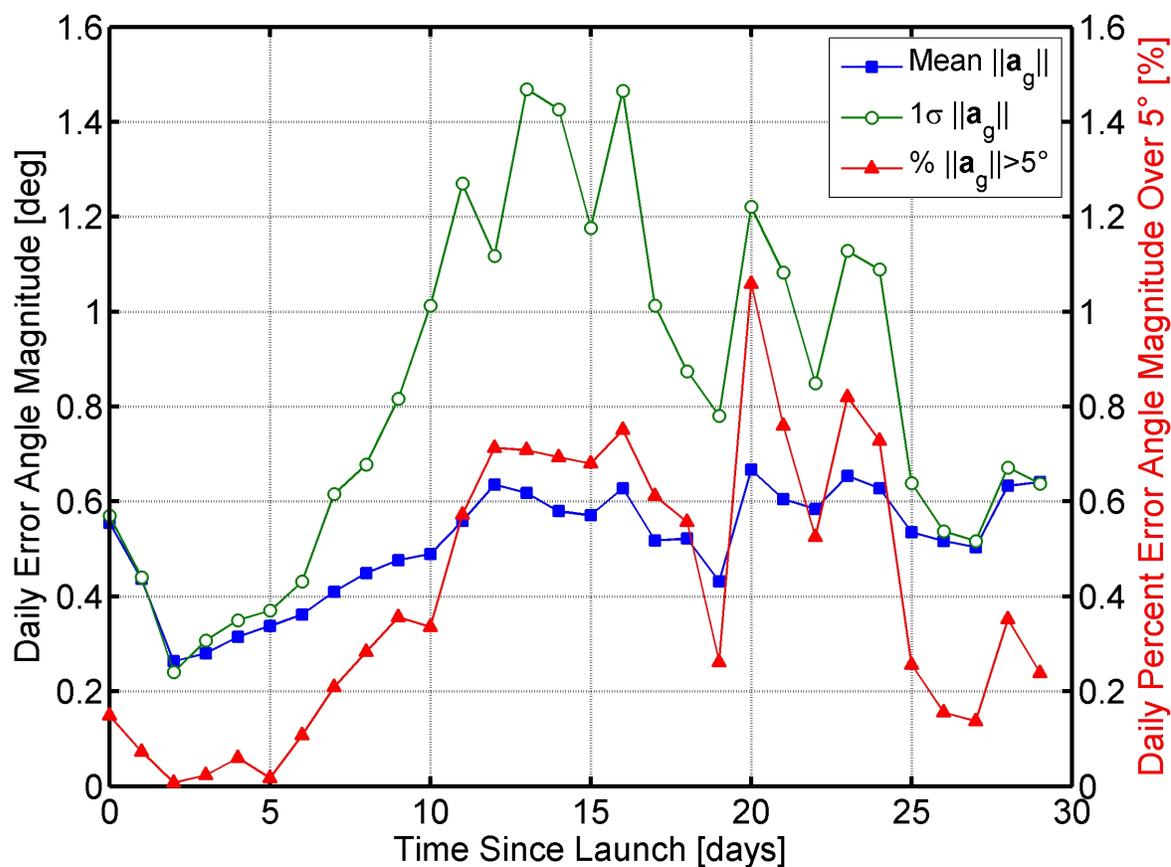
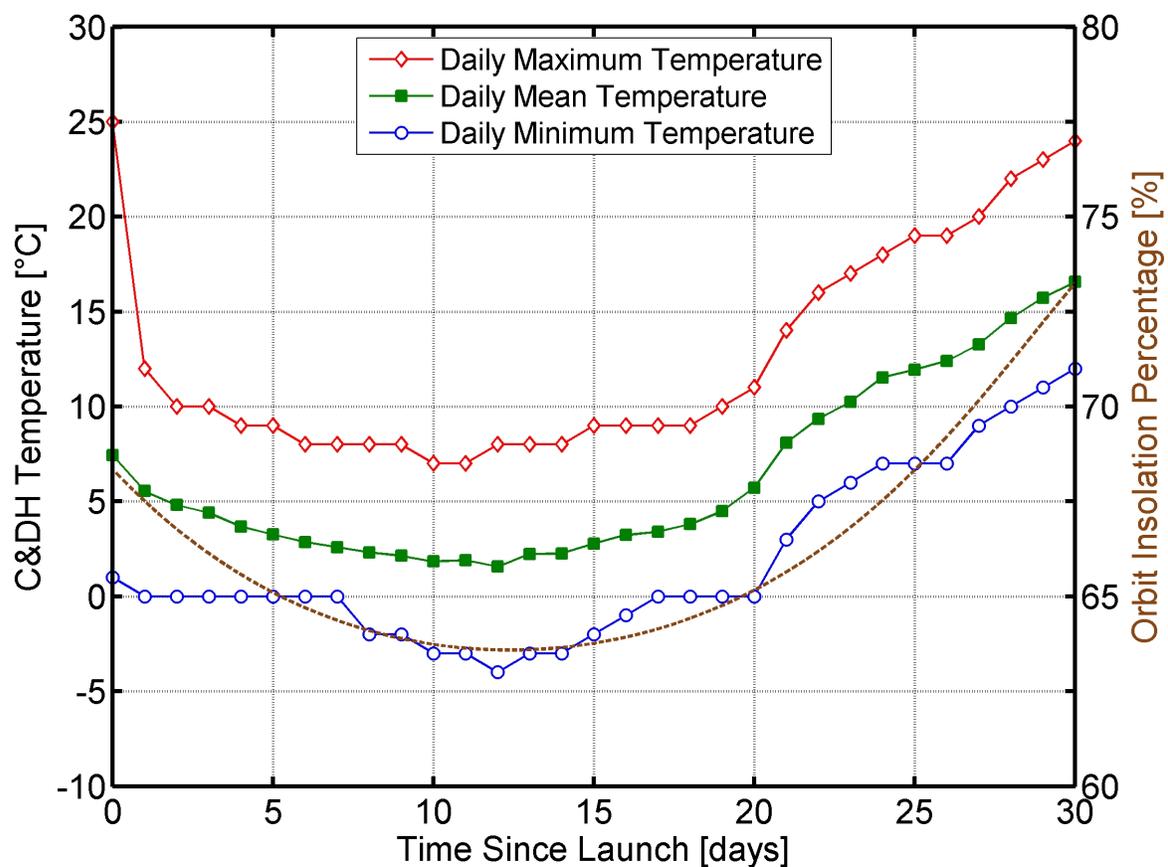


Figure 6.15: The daily maximum, mean, and minimum temperatures of the Command and Data Handling (C&DH) board (located within CSSWE) are represented by the red diamonds, green squares, and blue circles, respectively. The percentage of time the satellite is in the sun each day is represented by the brown dashed line.



increase during the eclipse period. In short, the MEKF is performing well.

### 6.3.1.3 Post-settling On-orbit Performance

Figure 6.18 and 6.19 show the  $\beta$  angle and angular velocity components with their associated  $3\sigma$  bounds for a 100 minute period about ten days after orbit insertion. The angular velocity of the X and Y axes varies from  $\pm 0.2^\circ/\text{s}$  and the roll rate is about  $-0.3^\circ/\text{s}$ . The roll rate estimate shows a minor discontinuity when the photodiodes enter use after the eclipse period. However, the discontinuity amplitude is within the uncertainty bounds of the fit, indicating that the filter is estimating uncertainty well.

CSSWE has now settled to within  $5^\circ$  of the  $\mathbf{B}$ -field, though the  $3\sigma$  uncertainty bounds of  $\beta$  sometimes reach large values, especially during eclipse. Note that the angular uncertainty is based on the assumption that the scaled error Gibbs vector uncertainty is a small angle and thus approximates the pitch, roll, and yaw uncertainties. Thus, uncertainties much greater than  $5^\circ$  should not be trusted. The uncertainty peaks during insolation (both  $\beta$  and angular velocity) are due to a temporary loss of the photodiodes due to the sun position in the body frame, as shown in Figure 6.20. The figure also shows the magnitude of the filter error angle; the small angle assumption is only broken once, during the transition from eclipse to insolation. Thus, even during the periods when the filter output uncertainty is suspect, the fitted attitude remains trustworthy.

### 6.3.2 Attitude Determination Validation

The CSSWE attitude determination results may be validated using independently-measured on-orbit data. Two events which validate the filtered attitude output are detailed below.

#### 6.3.2.1 Antenna Deployment Event

Due to requirements set by the CubeSat Design Specification [13], the satellite was launched with the antenna stowed within the CubeSat. CSSWE was programmed to automatically deploy the antenna two hours after orbit insertion; the satellite event log vouches for the deployment. The

Figure 6.16: The angle between the local magnetic field and the  $+^{\mathcal{B}}Z$  axis over a 100 minute period five hours after orbit insertion. The  $3\sigma$  uncertainty bounds are shown in red and the insolation times are marked in blue.

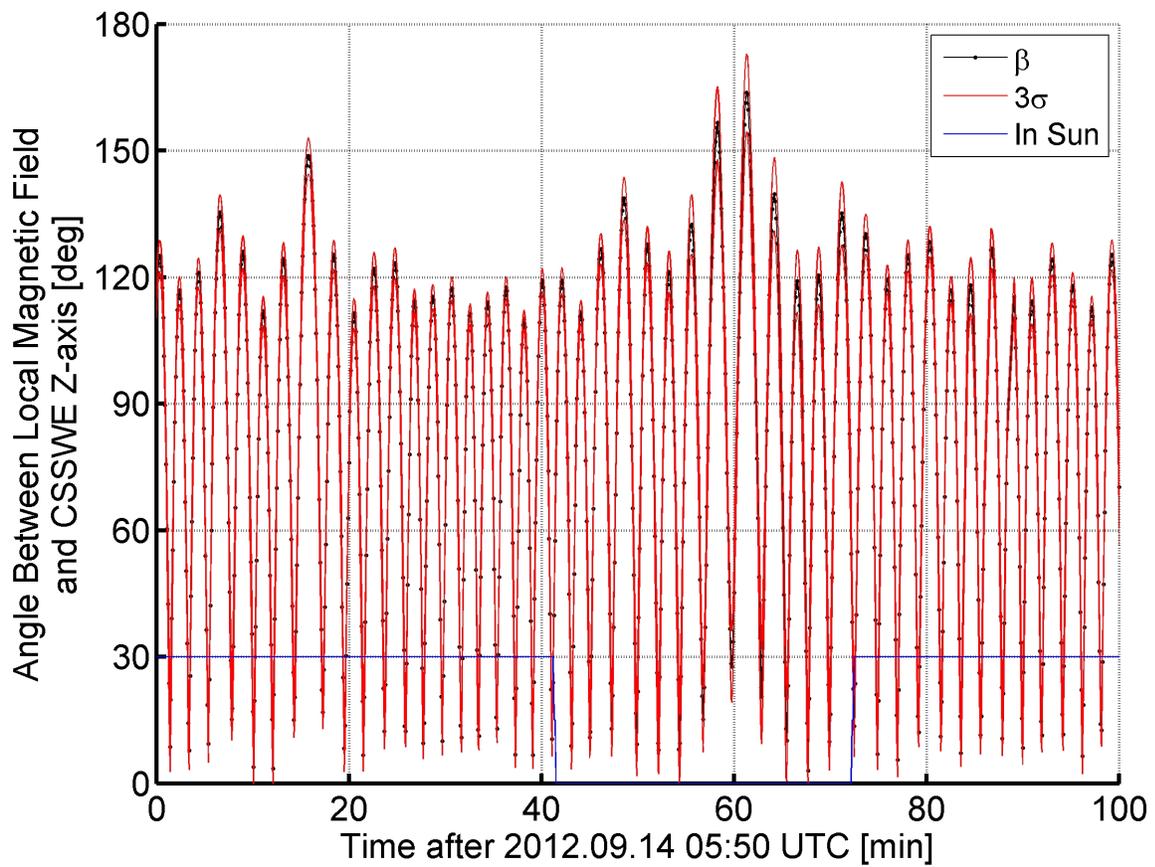


Figure 6.17: The satellite angular velocity components over a 100 minute period five hours after orbit insertion. The  $3\sigma$  uncertainty bounds are shown in red and the insolation times are marked in blue.

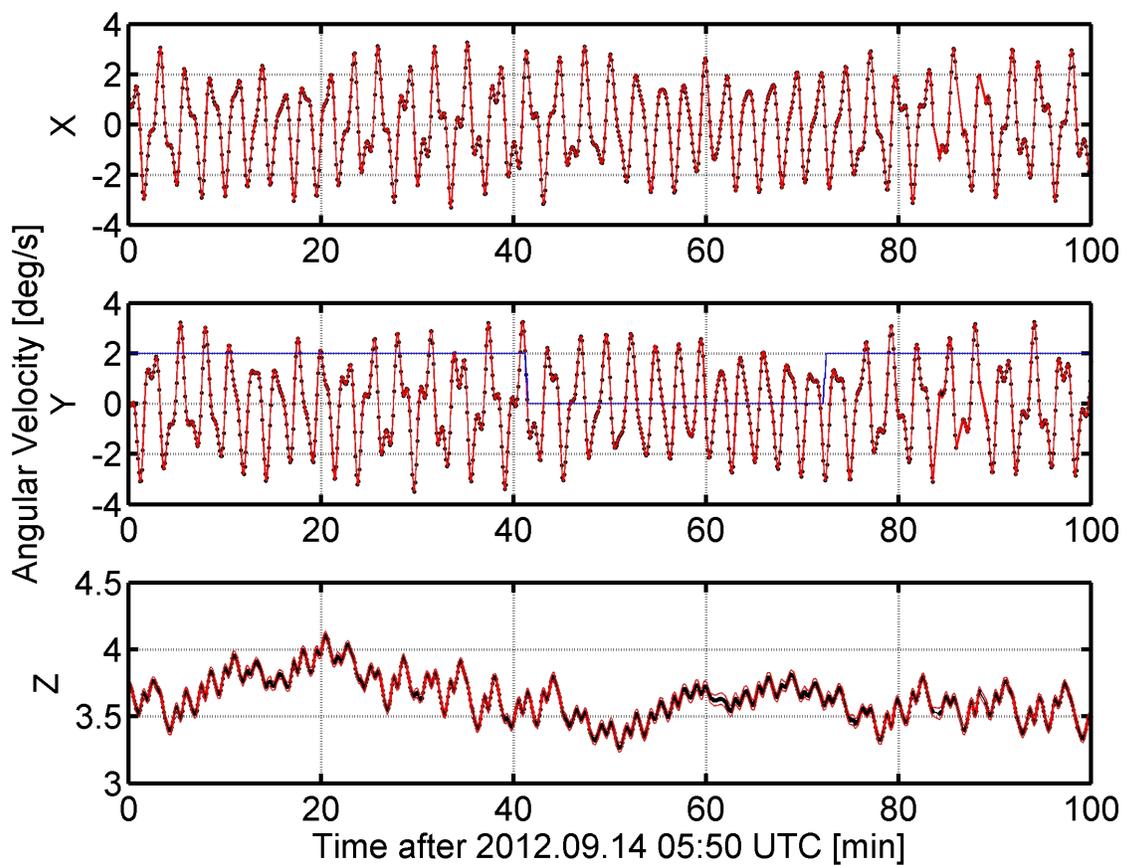


Figure 6.18: The angle between the local magnetic field and the  $+^{\mathcal{B}}Z$  axis over a 100 minute period ten days after launch. The  $3\sigma$  uncertainty bounds are shown in red and the insolation times are marked in blue.

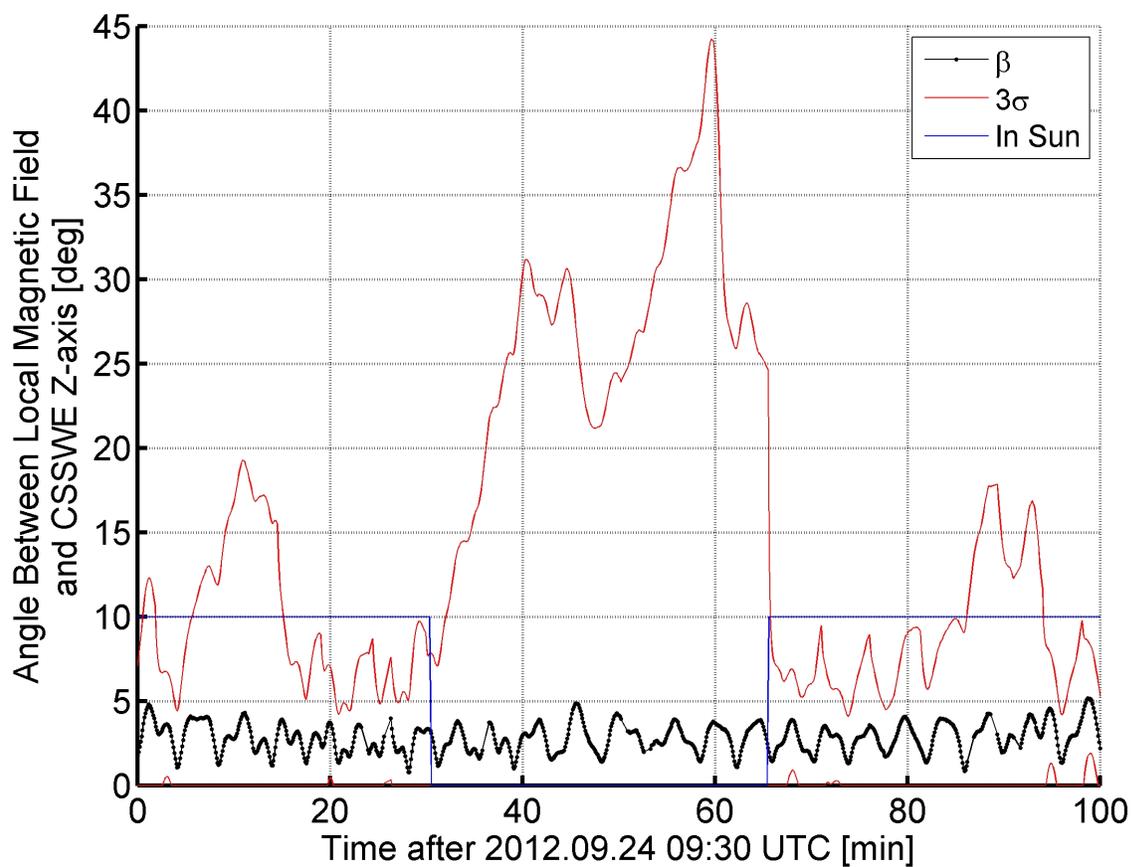


Figure 6.19: The satellite angular velocity components over a 100 minute period ten days after launch. The  $3\sigma$  uncertainty bounds are shown in red and the insolation times are marked in blue.

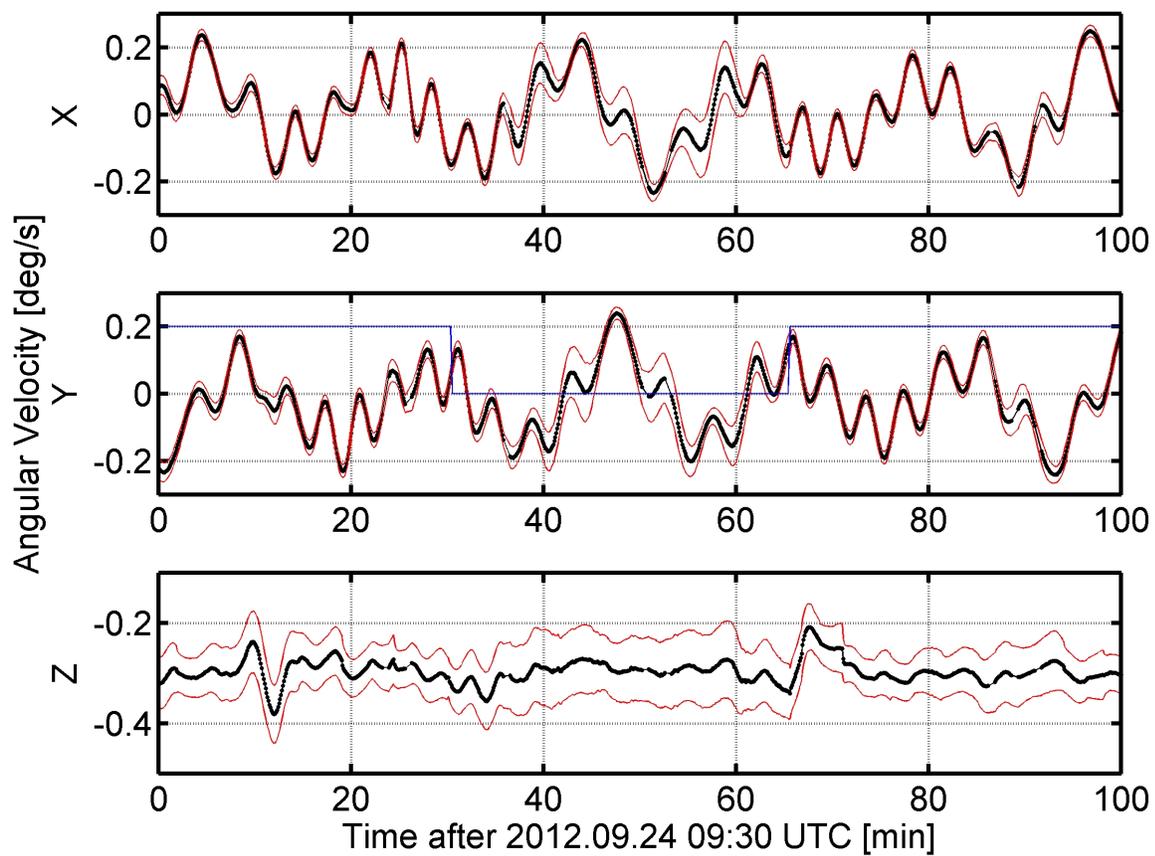
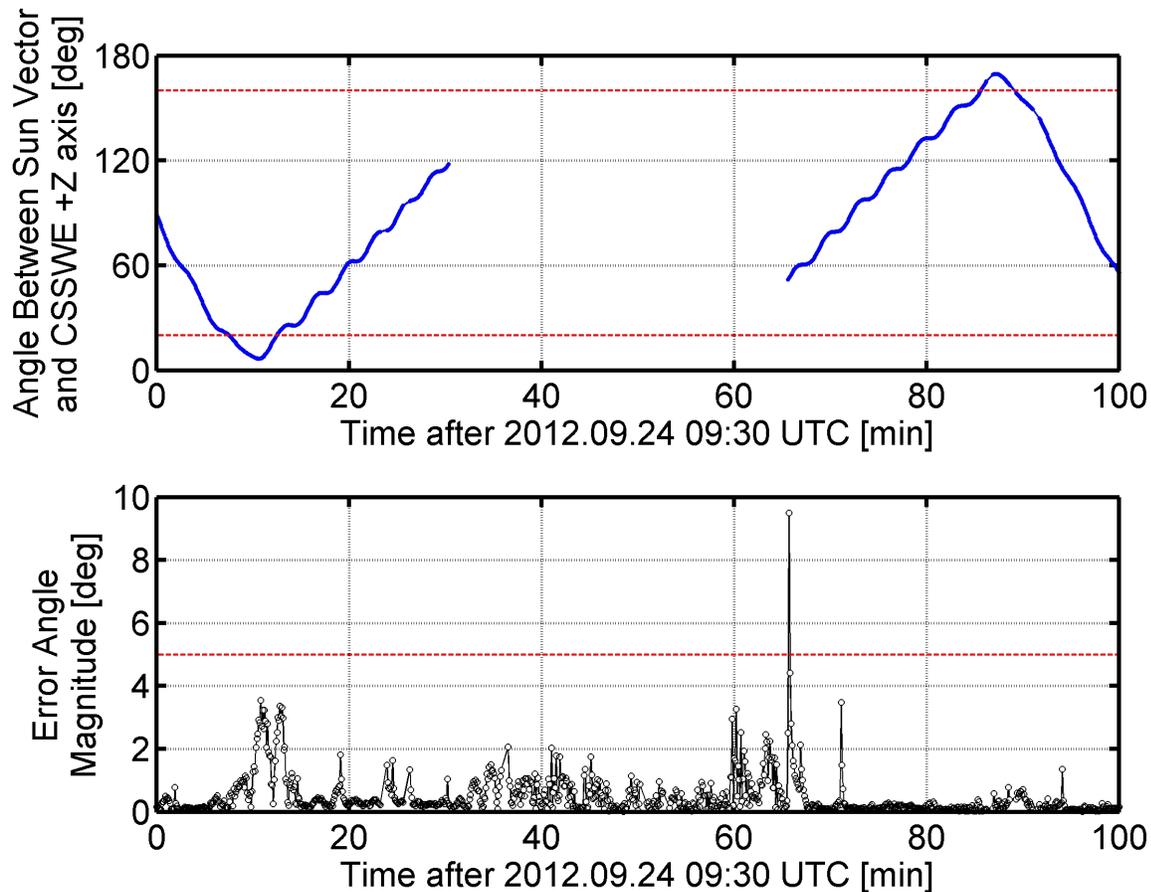


Figure 6.20: The top plot shows the angle between the sun vector and the CSSWE  $+^B Z$  axis (blue) and the high incidence angle cutoffs at  $20^\circ$  and  $160^\circ$  (dotted red lines). The uncertainty peaks during insolated times occur when the sun vector is within  $20^\circ$  of  $\pm^B Z$ , suggesting that the loss of the photodiodes causes the rapid increase in uncertainty. The bottom plot shows the magnitude of the scaled error Gibbs vector over the first 100 minutes (black circles) with the  $5^\circ$  limit to the small angle assumption denoted (red dotted line). Only the point directly following the transition back to insolation has an angular magnitude greater than five degrees.



spring-steel antenna deployment was expected to slightly change the magnetometer calibration parameters because of its ferromagnetic makeup. Instead, analysis of the on-orbit data shows that the deployment of the spring-steel antenna significantly changes both the magnetometer calibration and the magnetic moment of the satellite. The satellite motion difference was such that two different satellite magnetic moments were calculated during the fitting process described in Section 6.2.2:  $0.84 \text{ A}\cdot\text{m}^2$  and  $0.55 \text{ A}\cdot\text{m}^2$  before and after the antenna deployment, respectively. Figure 6.21 shows the system energy before and after the deployment event (denoted with the green line at 2 hours after orbit insertion); there is a significant decrease in the system rotational energy when the antenna deploys. Both the potential and kinetic rotational energy dynamics change significantly after the antenna deployment because the magnetic torque of the satellite dominates all other external torques. The antenna deployment does not create a visible difference in the filtered attitude or angular velocity estimates, indicating that the torque due to antenna deployment is insignificant.

### **6.3.2.2 Solar Panel Temperature Distribution**

The satellite roll rate has many implications. Much like meat roasting on a spit, when the roll rate approaches zero one side of the satellite can burn while the opposite side freezes. Figure 6.22 shows that the solar panel temperatures can reach extreme values when the roll rate is near zero. The effect is so pronounced that the panels can reach nearly  $100^\circ\text{C}$  even during a period of maximum eclipse time. The independently-measured solar panel temperature data are in good agreement with the MEKF fit.

Figure 6.21: Satellite rotational energy before and after the antenna deployment event (represented by the green line). The top plot shows the kinetic (red) and potential (blue) rotational energy while the bottom plot shows the total rotational energy.

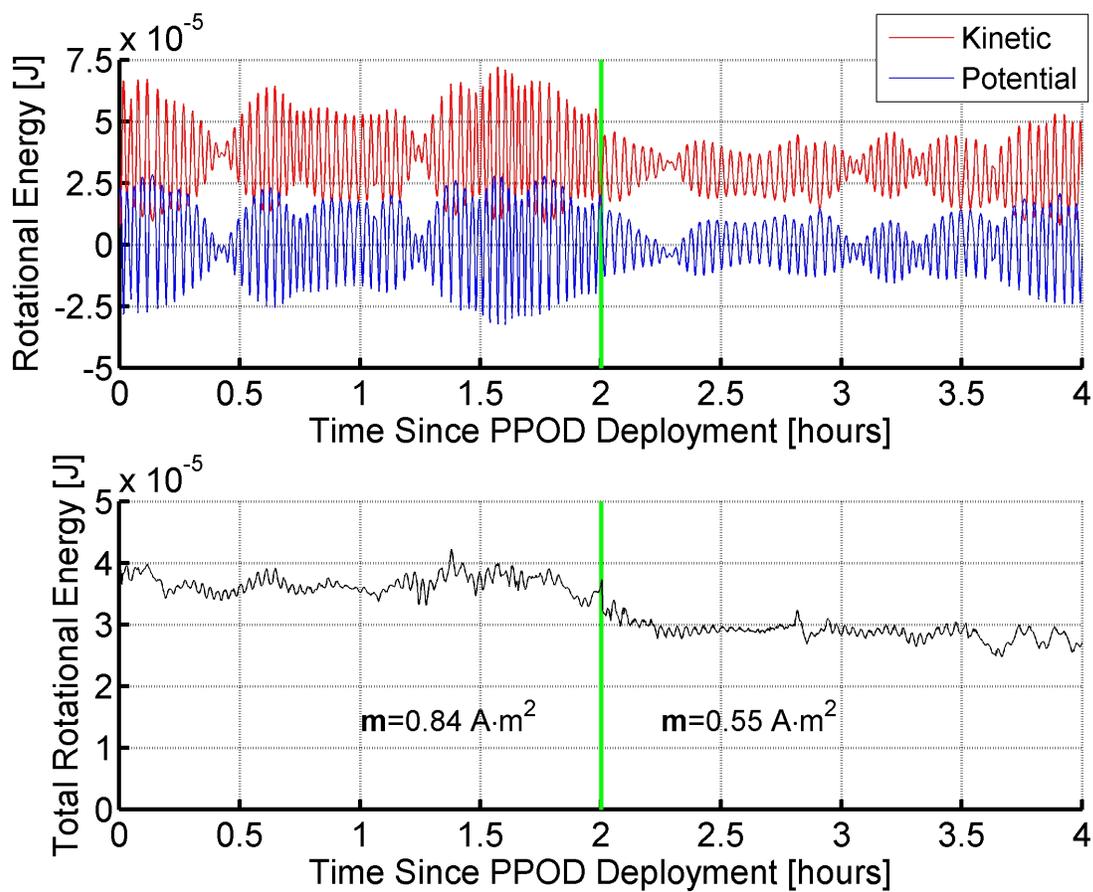


Figure 6.22: The top plot shows the maximum (open diamond), mean (filled square), and minimum (open circle) temperature of each of four solar panels; the orbit insolation percentage is overlaid (brown dotted line). The bottom plot shows the absolute value of the roll rate as estimated by the MEKF.

