Three-Axis Attitude Determination via Kalman Filtering of Magnetometer Data

by François Martel[†], Parimal K. Pal[*], and Mark L. Psiaki[**]

Abstract

A three-axis Magnetometer/Kalman Filter attitude determination system for a spacecraft in low-altitude Earth orbit is developed, analyzed, and simulation tested. The motivation for developing this system is to achieve light weight and low cost for an attitude determination system.

The extended Kalman filter estimates the attitude, attitude rates, and constant disturbance torques. Accuracy near that of the International Geomagnetic Reference Field model is achieved. Covariance computation and simulation testing demonstrate the filter's accuracy. One test case, a gravity-gradient stabilized spacecraft with a pitch momentum wheel and a magnetically-anchored damper, is a real satellite on which this attitude determination system will be used.

This work is similar to that of Heyler [5]. The application to a nadir pointing satellite and the estimation of disturbance torques represent the significant extensions contributed by this paper. Beyond its usefulness purely for attitude determination, this system could be used as a part of a low-cost three-axis attitude stabilization system.

Three-Axis Attitude Determination via Kalman Filtering of Magnetometer Data

by François Martel, Parimal K. Pal, and Mark L. Psiaki

1 Introduction

1.1 Objective

The objective of this work has been to develop a low-cost system for estimation of 3-axis spacecraft attitude information based solely on 3-axis magnetometer measurements from one satellite orbit. Such a system will be useful for missions that operate in an inclined, low-Earth orbit and require only coarse attitude information. It can also serve as the sensor part of a low-cost 3-axis closed-loop attitude control system, or as a back-up attitude estimator.

A single 3-axis magnetometer measurement can give only 2-axes worth of attitude information and no attitude rate or disturbance torque information. Therefore, this attitude determination system must use a sequence of magnetometer measurements. It processes these measurements recursively in a Kalman filter. This paper, then, describes the design, development, analysis, and simulation testing of a Kalman filter and reports its expected performance. A follow-on, post-launch paper is planned to report actual performance.

1.2 Background/Prior Work

Kalman filters have been widely applied to the problem of spacecraft attitude determination [1-7]. Everything from star sensors [2,3] to sun sensors [4], gyroscopes [2], and magnetometers [4,5] have been used for filter inputs, and accuracies as fine as 2 arc sec. are possible [3].

Very few attitude determination systems have attempted to use only magnetometer data to estimate attitude. Perhaps this is because of the low accuracy of the measurements; even with perfect magnetometer measurements, inaccuracy of the knowledge of the Earth's magnetic field may introduce errors of 0.4° per axis. Perhaps such systems are rare because of the complexity of computing the Earth's magnetic field from spherical harmonic models [6]. In at least one case the benefits (low cost and low weight) have outweighed the costs and such a system has been developed. Heyler reports the use of such a system on the NOVA program [5]. That system was able to estimate spin axis attitude with a 2° accuracy as well as spin rate. These estimates were based on one eighth of an orbit's worth of magnetometer readings.

The Kalman filter reported in this paper uses 50 to 300 magnetometer samples distributed evenly over an orbit to estimate 3-axis attitude, attitude rate, and disturbance torques for a gravity-gradient-stabilized spacecraft. It is similar to the filter described by Heyler in that 3-axis information is derived purely from magnetometer measurement time histories. It differs from Heyler's filter in two respects; it estimates the attitude and rates for a different type of spacecraft, and it estimates disturbance torques. Also presented is a detailed accounting of the various contributions to estimation error, including the effects of spacecraft dynamic modeling error.

1.3 Outline of Approach

The remainder of this paper contains descriptions of the dynamic model of the spacecraft under consideration, the filter design, and the filter evaluation criteria and procedures. It concludes with the results of the filter evaluation. The spacecraft description discusses the type of spacecraft for which this filter will work and presents notation and equations necessary to the remaining sections. The filter design section presents the overall filter structure and two different gain selection techniques. The section on evaluation methodology describes the filter accuracy and stability performance criteria and the tools that were used to gauge these properties. The results of the accuracy and stability evaluations are presented in the final section, which includes examples of simulation time histories as well as numerical measures of performance.

2 Spacecraft Dynamic Model

2.1 Mission/Orbit Characteristics

The Kalman filter discussed in this work is applicable to nadir pointing Earth satellites operating at low altitudes in inclined orbits. The inclination and low altitude of the orbit are necessary to the proper functioning of the filter. The orbit must stay close enough to the Earth, within about 4 Earth radii [6], so that a spherical harmonic approximation of the Earth's magnetic field gives a reliable attitude reference. Some inclination of the orbit is necessary to make the attitude of all three axes sufficiently observable. Pitch information in a 1-orbit magnetometer time history gets poor for low inclinations, although theoretically, there is still some pitch information even in equatorial orbits; the Earth's magnetic poles do not coincide with its rotational poles. This study considers spacecraft in nearly circular orbits at 1.1 to 1.2 Earth radii. Filter analysis and testing has been done for the inclinations 43° and 57°.

2.2 Spacecraft Attitude Dynamics Model

The generic spacecraft (S/C) under consideration is a gravity gradient stabilized spacecraft. One model also has a pitch momentum wheel for passive yaw stiffening and a magnetically anchored damper for passive libration damping. The following equations of motion model the spacecraft attitude dynamics for purposes of filter state propagation:

$$\dot{\omega} = I_{inrt}^{-1}[n - \omega \times (I_{inrt}\omega + h_w)] \tag{1}$$

$$\dot{q} = \frac{1}{2}\begin{bmatrix} 0 & \omega_{SC/E3} & -\omega_{SC/E2} & \omega_{SC/E1} \\ -\omega_{SC/E3} & 0 & \omega_{SC/E1} & \omega_{SC/E2} \\ \omega_{SC/E2} & -\omega_{SC/E1} & 0 & \omega_{SC/E3} \\ -\omega_{SC/E1} & -\omega_{SC/E2} & -\omega_{SC/E3} & 0 \end{bmatrix} q \tag{2}$$

$$\dot{n}_d = 0 \tag{3}$$

where ω is the S/C's inertial angular velocity vector, I_{inrt} is the moment and product of inertia matrix, n is the total external vector torque acting on the S/C, h_w is the constant vector angular momentum of the pitch wheel, q is a quaternion that represents the orientation of the S/C-fixed coordinate system with respect to an Earth-fixed coordinate system, $\Omega_{SC/E}$ is the S/C's Earth-relative angular velocity,

and n_d is the disturbance torque (the net unmodeled external torque). All of the above are expressed in S/C-fixed coordinates except the quaternion. It is expressed in Earth-fixed coordinates. Equation 1 is Euler's equation for rigid body rotational dynamics, and eq. 2 is the kinematic equation for a quaternion [6]. Equation 3 is special to the filter. It represents the unmodeled disturbance torques.

The net external torque acting on the S/C, n, has been divided into three components, gravity gradient torque, n_{gg}, passive magnetically-anchored damper torque, n_{damp}, and all other unmodeled disturbance torques, n_d:

$$n = n_{gg} + n_{damp} + n_d \tag{4}$$

The first two of these torque components, when present, have been explicitly modeled for purposes of filter state and covariance propagation.

The gravity gradient torque depends on the attitude quaternion, the ephemeris, and the moments and products of inertia:

$$n_{gg} = n_{gg}(q, t; I_{inrt}) \tag{5}$$

where t is the time. The gravity gradient model used in this study neglects J_2 effects [6].

The magnetically-anchored damper torque depends on the S/C-fixed magnetic field unit vector and its time rate of change, which in turn, depend on the attitude quaternion, the Earth-relative S/C angular velocity, and the ephemeris [6]:

$$n_{damp} = c_{damp}\left[\hat{b}(q, t) \times \frac{d\hat{b}}{dt}\bigg|(q, \Omega_{SC/E}, t) \right] \tag{6}$$

where c_{damp} is the damping factor, $\hat{b}$ is the magnetic field unit vector in S/C-fixed coordinates, and the derivative with respect to time is the total derivative (q is time varying).

The unmodeled disturbance torque, n_d, may include the effects of atmospheric drag, solar radiation pressure, residual magnetic dipole moment, S/C dynamics modeling errors, or any other unmodeled external torques. No explicit physical model of any of these torques is included. Rather, this term is retained in an effort to estimate these torques in the filter by modeling them as a random walk process.

The coordinate systems used in this study are a S/C-fixed coordinate system, an Earth-fixed coordinate system,

Table 1

Attitude Dynamics Parameters of Two Spacecraft

Spacecraft No.	I_{xx}	I_{yy}	I_{zz}	I_{xy}	I_{xz}	I_{yz}	c_{damp}	h_{w_2}
	----------------------- (kg-m^2) -------------------						(N-m-s)	(kg-m^2/s)
1	200,000	300,000	70,000	200	2,000	-60	1	-70
2	250	250	10	0	0	0	0	0

and an orbit-following coordinate system. The S/C-fixed coordinate system is a Roll-Pitch-Yaw coordinate system; the x axis is nominally* parallel to the velocity vector, the y axis is nominally anti-parallel to the orbit normal, and the z axis is nominally along nadir. This reference frame is used to define the equations of motion and related equations, eq. 1-6, the inertia matrix, I_{inrt}, and the pitch wheel angular momentum, h_w.

The orbit-following coordinate system defines the nominal orientation of the gravity-gradient-stabilized S/C. Its z axis is exactly along nadir, its y axis is exactly anti-parallel to orbit normal, and its x axis is approximately parallel to velocity (exactly parallel in the case of circular, nondecaying orbits). Its only purpose in this study is as a point of reference for measuring roll, pitch, and yaw angles in reporting attitude results.

The Earth-fixed coordinate system has its origin at the Earth's center. Its x axis passes through the equator at the Greenwich meridian, its y axis passes through the Equator at 90° East Longitude, and its z axis passes through the North Pole. It is used to calculate the S/C ephemeris and the Earth's magnetic field, which are used in torque modeling and filter update calculations. Because this reference frame rotates with the Earth, there is a difference between the S/C's inertial angular velocity, ω, and its angular velocity with respect to this reference frame, $\omega_{SC/E}$:

$$\omega_{SC/E} = \omega - A \begin{bmatrix} 0 \\ 0 \\ \omega_e \end{bmatrix} \tag{7}$$

where A is the coordinate transformation matrix from Earth-fixed to S/C-fixed coordinates defined by q, and $\omega_e = 7.29 \times 10^{-5}$ rad/sec is the Earth's rotational angular velocity. The angular velocity of the Earth as its revolves about the Sun has been neglected in this transformation.

Table 1 lists the nominal values of the attitude dynamics parameters for two S/C examples. Spacecraft 1 is stabilized by a long gravity gradient boom with a tip mass, a constant momentum pitch wheel, and a magnetically-anchored damper. Spacecraft 2 has a gravity-gradient boom, but it is left neutrally stable in yaw. The tabulated parameter values (sometimes with deliberately introduced perturbations) apply to the analyses and simulations described below.

2.3 Attitude Determination Hardware

The only attitude determination sensor used by this filter is a 3-axis magnetometer. It measures the magnetic field vector in S/C-fixed coordinates:

$$b = A \, b_{E/F} \tag{8}$$

where $b_{E/F}$, the magnetic field in the Earth-Fixed coordinate system, depends only on the S/C ephemeris. The A matrix depends on q, so eq. 8 defines the nonlinear measurement equation used by the extended Kalman filter.

---

* In the absence of orbital eccentricity, librational motion, disturbance torques, or product of inertia terms.

2.4 A Linearized Attitude Dynamics Model

Linearized equations of motion and sensor equations are useful for filter analysis and design. This involves linearization of eq. 1, 2, 4, 5, 6, and 8. They are linearized about the nominal S/C attitude time history: z axis along nadir, y axis along negative orbit normal, and y-axis angular velocity equal to the orbital rate. The orbit is assumed circular and J_2 effects are neglected. As a further simplification, a dipole model of the Earth's magnetic field is used [6], and the field at the S/C is assumed periodic with the orbital period (the rotation of the field with the Earth is ignored).

The attitude quaternion has been linearized in a special way. Instead if expressing q in terms of the sum of a nominal value plus a perturbation, it is expressed in terms of a perturbation quaternion times the nominal quaternion using quaternion multiplication:

$$q = q_{nom} \bullet \begin{bmatrix} \Delta q_1 \\ \Delta q_2 \\ \Delta q_3 \\ 1 \end{bmatrix} = q_{nom} \bullet \begin{bmatrix} \Delta q \\ 1 \end{bmatrix} \quad (9)$$

where, by definition of the nominal attitude time history, q_{nom} defines the attitude of the orbit-following coordinate system. The perturbational quaternion is already normalized to within first order in the Δq_i. This perturbational expression of the attitude has just three unknowns; the fourth is not needed because angles are small, the equations are linear, and no attitude singularity occurs.

The linearized equations are

$$\Delta \dot{\omega} = I_{inrt}^{-1}[\Delta n - \Delta \omega \times (I_{inrt}\omega_{orb} + h_w)$$

$$- \omega_{orb} \times (I_{inrt}\Delta \omega)] \quad (10)$$

$$\Delta \dot{q} = \frac{1}{2}\Delta \omega \quad (11)$$

$$\Delta n = \Delta n_{gg} + \Delta n_{damp} + n_d \quad (12)$$

$$\Delta n_{gg} = \frac{3\mu}{r_{S/C}^3}\left\{ \begin{bmatrix} -I_{23} \\ I_{13} \\ 0 \end{bmatrix} + 2\begin{bmatrix} I_{33}-I_{22} & I_{12} & 0 \\ I_{12} & I_{33}-I_{11} & 0 \\ -I_{13} & -I_{23} & 0 \end{bmatrix}\Delta q \right\}$$

$$(13)$$

$$\Delta n_{damp} = c_{damp}\left[(-I + \hat{b}_{orb}\hat{b}_{orb}^T)\Delta \omega + \hat{b}_{orb} \times \frac{d\hat{b}_{orb}}{dt} \right.$$

$$\left. + 2\left(\hat{b}_{orb} \times \frac{d\hat{b}_{orb}}{dt}\right) \times \Delta q \right] \quad (14)$$

$$b = b_{orb} - 2\Delta q \times b_{orb} \quad (15)$$

where $\Delta \omega$ is the perturbational S/C angular velocity expressed in S/C-fixed coordinates, ω_{orb} is the orbital angular velocity expressed in orbit-following coordinates (its only nonzero element is its y element), μ is the geocentric gravitational constant, $r_{S/C}$ is the S/C geocentric radius, I_{ij} is the i,j element of I_{inrt}, I is the identity matrix, and b_{orb} is the Earth's magnetic field vector at the S/C expressed in orbit-following coordinates.

These equations can be combined in standard state vector format to yield a 9th-order system of the form

$$\Delta \dot{x} = F(t)\Delta x + z(t) \quad (16)$$

$$y = H(t)\Delta x \quad (17)$$

where the state is defined as $\Delta x^T = (\Delta \omega^T, \Delta q^T, n_d^T)$ and where the observation is $y = \hat{b} \times \hat{b}_{orb}$. This definition of y retains all of the attitude information in the magnetometer measurements and gives an H(t) matrix consistent with the innovation definition given below (eq. 21). The 9x9 F(t) and matrix and the 9-element z(t) vector are derived from eq. 3, 10-14 and the definition of Δx. The 3x9 H(t) matrix is derived from eq. 15 and the definition of y. F(t), H(t), and z(t) are all periodic at the orbital period because the magnetic field has been assumed periodic at the orbital period. The periodicity of this linear system can be used to advantage in filter design and analysis.

The presence of z indicates that linearization has not been done about the nominal motion. As can be seen from eq. 13 and 14, the nonhomogeneous terms result from product of inertia terms (a gravity gradient effect) and from the time variation of the Earth's magnetic field as experienced in the orbit following reference frame (a magnetically-anchored damper effect). Nonzero z means that the S/C is not exactly trimmed at its nominal orientation. This out-of-trim condition is not vary large ($\leq \sim 1^\circ$), and the linearized model is a good approximation for small perturbations from trim.

3 Filter Design

3.1 Filter Mission

The filter's mission is to estimate 3-axis attitude, attitude rate, and disturbance torque. The accuracy goal of the attitude estimates is on the order of $\pm 1°$. This information may be required to run experiments off of a passively stabilized S/C or to provide feedback signals for active stabilization. For the former mission, the attitude estimation may be done in a ground station in batch mode once per orbit. When part of a feedback control loop, the filter will operate on board the S/C, recursively updating the attitude, rate, and torque estimates.

The filter computer program must execute quickly for such missions. When operating in a ground station there is only a short time window for magnetometer data transmission, one orbit's worth of filtering, and subsequent experimentation. Less time spent filtering leaves more time for the primary mission experiments. When operating on board, the filter has more time to filter one orbit's worth of data, but filtering will be only one of many tasks for the on-board computer. Less time spent filtering means more computer time left for primary mission usage.

3.2 Filter Structure and Gain Computation

The basic filter structure is that of the typical sampled-data extended Kalman Filter: a state/covariance propagation phase alternating with a state/covariance update phase once for each sensor sample time. Figure 1 gives a block diagram of this basic structure and the associated information flow. In the figure, t_k and t_{k+1} are sample times, b_{meas} is the vector of magnetometer measurements, $\tilde{x}$ is the state estimate, P is the state estimate covariance matrix (not always used), and the (-) and (+) superscripts on $\tilde{x}$ and P refer to pre- and post-update values, respectively, at a given magnetometer sample instant.

The state propagation portion of the filter is the usual nonlinear simulation of the system equations of motion, eq. 1-3. Therefore, the state estimate in the extended Kalman filter is a 10-dimensional vector:

$$\tilde{x} = \begin{bmatrix} \tilde{\omega} \\ \tilde{q} \\ \tilde{n}_d \end{bmatrix} \tag{18}$$

where the (~) overstrike indicates an estimate. The state propagation algorithm computes $\tilde{x}^-(t_{k+1})$ as a function of $\tilde{x}^+(t_k)$ by numerically integrating eq. 1-3 from time t_k to time t_{k+1} starting from the initial conditions:

$$\begin{bmatrix} \tilde{\omega}(t_k) \\ \tilde{q}(t_k) \\ \tilde{n}_d(t_k) \end{bmatrix} = \tilde{x}^+(t_k) \tag{19}$$

Formally, one may consider this procedure the definition of a vector function, f, and a discrete-time system:

$$\tilde{x}^-(t_{k+1}) = f[\tilde{x}^+(t_k),k] \tag{20}$$

Filter state propagation and evaluation of the function f are equivalent.

The state update calculation in this filter is slightly different from the traditional extended Kalman filter update in several respects. The filter innovations, the method of updating the quaternion estimate, and the method of calculating the filter gain are all slightly different from standard extended Kalman filter practice. Each of these differences has been introduced in order to handle the nonlinearities in a manner better than brute force linearization.

The cross product of the measured magnetic field unit vector with its pre-update estimate has been chosen for the innovation:

$$\underline{v}(t_{k+1}) = \hat{b}_{meas}(t_{k+1}) \times \hat{\tilde{b}}^-(t_{k+1}) \tag{21}$$

where $\underline{v}$ is the innovation and where the (^) overstrike indicates a unit vector. The standard extended Kalman filter would simply take the difference between b_{meas} and its pre-update estimate to form the innovation. The formula in eq. 21 essentially throws out all of the length information in the measured magnetic field. Nothing is lost as there is no attitude information in the length. In the linear analysis, either innovation formula would give the same update, but the eq. 21 formula is to be preferred in the nonlinear case because its magnitude and direction both are physically significant; they define the magnitude and direction of the known angular error.

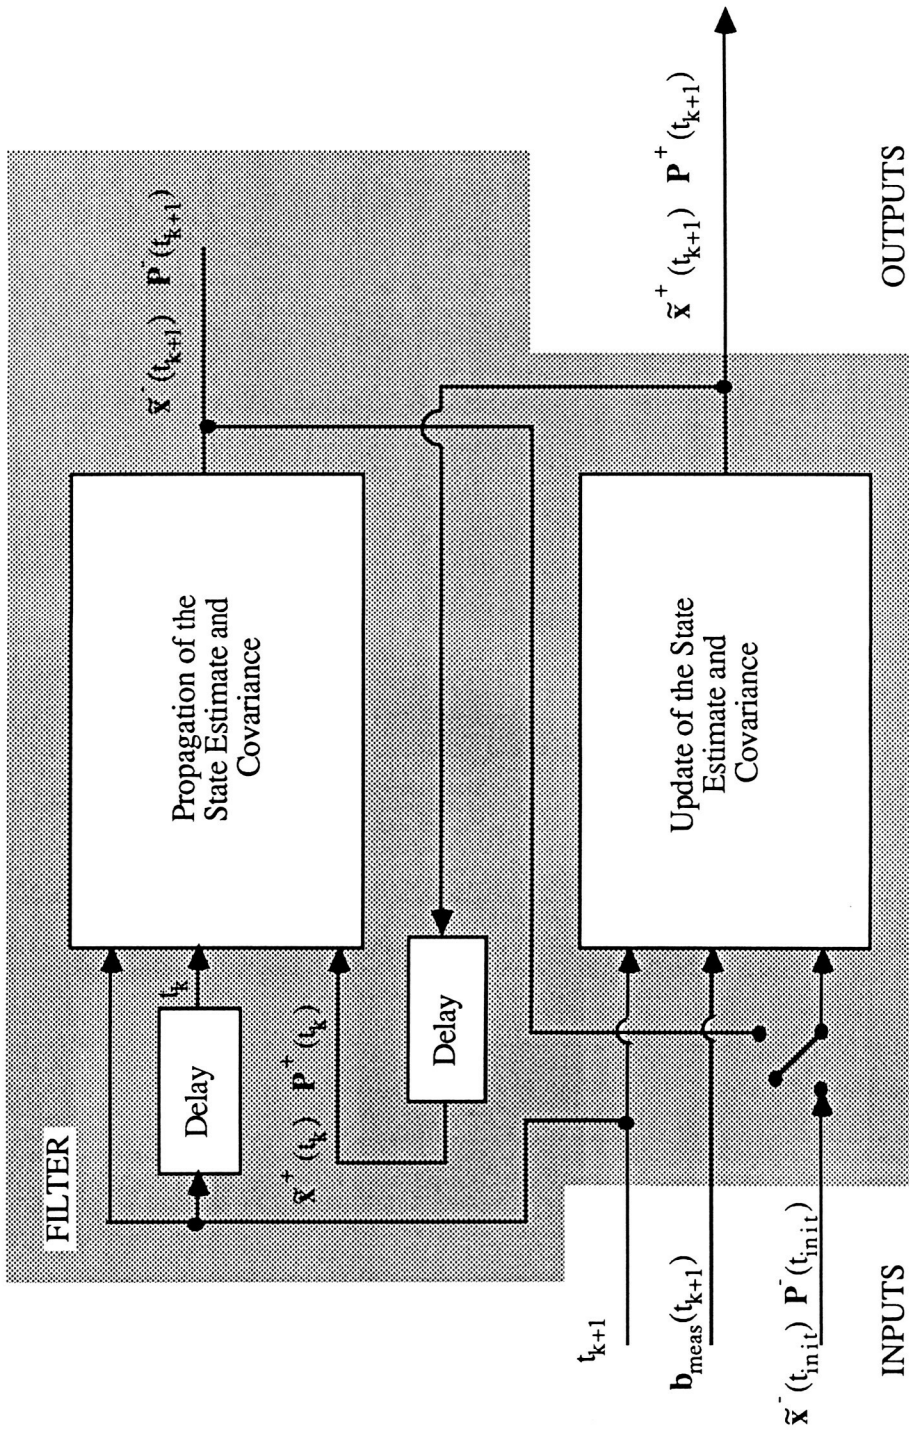The update formulas for the attitude rate and disturbance torque estimates take the usual form:

349

Figure 1. Block Diagram of Kalman Filter Structure and Information Flow.

FILTER

Propagation of the State Estimate and Covariance

Update of the State Estimate and Covariance

Delay

Delay

t_k

$\tilde{\mathbf{x}}^+(t_k)$ $\mathbf{P}^+(t_k)$

$\tilde{\mathbf{x}}^-(t_{k+1})$ $\mathbf{P}^-(t_{k+1})$

$\tilde{\mathbf{x}}^+(t_{k+1})$ $\mathbf{P}^+(t_{k+1})$

OUTPUTS

t_{k+1}

$\mathbf{b}_{meas}(t_{k+1})$

$\tilde{\mathbf{x}}^-(t_{init})$ $\mathbf{P}^-(t_{init})$

INPUTS

350

$$\tilde{\underline{\omega}}^+(t_{k+1}) = \tilde{\underline{\omega}}^-(t_{k+1}) + K_{\omega_{k+1}}\underline{v} \qquad (22)$$

$$\tilde{n}_d^+(t_{k+1}) = \tilde{n}_d^-(t_{k+1}) + K_{n_{k+1}}\underline{v} \qquad (23)$$

but the quaternion update uses quaternion multiplication instead of addition:

$$\tilde{q}^+(t_{k+1}) = \tilde{q}^-(t_{k+1}) \cdot \begin{bmatrix} \Delta q_{ud} \\ \sqrt{1 - |\Delta q_{ud}|^2} \end{bmatrix} \qquad (24)$$

with

$$\Delta q_{ud} = K_{q_{k+1}}\underline{v} \qquad (25)$$

This form of the quaternion update explicitly recognizes that there are only three free variables in the quaternion and updates it accordingly. The Δq_{ud} vector elements constitute the three free variables as in the quaternion update method described by Lefferts et.al. [1]. The update preserves the quaternion's normalization.

The three K matrices in eq 22, 23, and 25 are the Kalman filter gain matrices. Because of the form of the innovation and the attitude update, the magnitude of the K_q gain matrix takes on physical significance. If the attitude determination system attempted to eliminate all of the measured attitude error at each measurement, K_q would be 1/2 times the 3x3 identity matrix. This is because $\underline{v}$ and Δq_{ud} are both in the direction of the measured attitude error, but the magnitude of $\underline{v}$ is proportional to the sine of the error while the magnitude of Δq_{ud} is proportional to the sine of half the error.

Two Kalman filter gain selection schemes have been tried. One uses, with a few necessary modifications, the standard extended Kalman filter covariance propagation, gain calculation, and covariance update formulas. The other uses fixed gains. The traditional extended Kalman filter was used in order to get the best possible filter performance. The fixed gain filter was used in the hope of achieving acceptable performance at a greatly reduced computational load.

The extended Kalman filter gain and covariance equations perform the necessary bookkeeping to derive 3-axis attitude information from the sequence of single-vector observations available to this filter. Modifications to these equations are necessitated by the nonstandard quaternion update, eq. 24 and 25, and by the nonstandard innovation, eq. 21. In addition, a discrete-time form of the covariance

propagation equation is used in the interest of reducing the programming complexity; the continuous-time update would involve programming covariance differential equations; the discrete-time update involves only matrix arithmetic and calculation of the state transition matrix via numerical differentiation of the $f[\;]$ function (eq. 20).

The discrete-time covariance propagation equation is

$$P^-(t_{k+1}) = \Phi(t_{k+1},t_k)P^+(t_k)\Phi^T(t_{k+1},t_k) + Q \qquad (26)$$

where P is the 9x9 covariance matrix for the perturbation vector $\Delta x^T = (\Delta\underline{\omega}^T, \Delta q^T, \Delta n_d^T)$, $\Phi(t_{k+1},t_k)$ is the state transition from $\Delta x(t_k)$ to $\Delta x(t_{k+1})$, and Q is the discrete-time disturbance input covariance matrix. The Δq perturbation is defined in the same way as the Δq_{ud} perturbation in eq. 24. This use of the 3-element Δq instead of the 4-element Δq simplifies the filter because there is no normalization constraint on the elements of Δq. Thus, the 9x9 covariance is nonsingular, and the standard Kalman filter equations for gain computation and covariance update can be used [1].

The use of Δq affects the computation of the state transition matrix. Normally, it would simply be the 10x10 matrix $\partial f/\partial x$. The expression of quaternion perturbations in terms of three independent components makes the state transition matrix

$$\Phi(t_{k+1},t_k) =$$

$$\begin{bmatrix} I & 0 & 0 \\ 0 & \frac{\partial q}{\partial \Delta q}\Big|^T_{\tilde{q}^-(t_{k+1})} & 0 \\ 0 & 0 & I \end{bmatrix} \frac{\partial f}{\partial x}\Big|_{[\tilde{x}^+(t_k),k]} \begin{bmatrix} I & 0 & 0 \\ 0 & \frac{\partial q}{\partial \Delta q}\Big|_{\tilde{q}^+(t_k)} & 0 \\ 0 & 0 & I \end{bmatrix}$$

$$(27)$$

which is a 9x9 matrix. The derivative of the quaternion with respect Δq is just

$$\frac{\partial q}{\partial \Delta q}\Big|_{\tilde{q}} = \begin{bmatrix} \tilde{q}_4 & -\tilde{q}_3 & \tilde{q}_2 \\ \tilde{q}_3 & \tilde{q}_4 & -\tilde{q}_1 \\ -\tilde{q}_2 & \tilde{q}_1 & \tilde{q}_4 \\ -\tilde{q}_1 & -\tilde{q}_2 & -\tilde{q}_3 \end{bmatrix} \qquad (28)$$

where the $\tilde{q}_i$ are the elements of $\tilde{q}$. This completes the covariance propagation formula. The approach of Lefferts et. al. is essentially the same [1].

351

The gain computation and covariance update formulæ are the usual Kalman filter formulæ [8]:

$$
\begin{bmatrix} K_{\omega_{k+1}} \\ K_{q_{k+1}} \\ K_{n_{k+1}} \end{bmatrix} = K_{k+1}
$$

$$
= P^-(t_{k+1})H^T(t_{k+1})\left[R + H(t_{k+1})P^-(t_{k+1})H^T(t_{k+1})\right]^{-1}
$$

(29)

$$
P^+(t_{k+1}) = P^-(t_{k+1}) - K_{k+1}H(t_{k+1})P^-(t_{k+1})
$$

(30)

where R is the measurement noise covariance matrix and K_{k+1} is the 9x3 filter gain matrix. The calculation of the observation matrix, $H(t_{k+1})$ (as in eq. 17), accounts for the nonstandard innovation and the nonstandard quaternion update. The following formula gives the true H in the spirit of the extended Kalman filter:

$$
H(t_{k+1}) = -\frac{\partial \underline{v}(t_{k+1})}{\partial \Delta x(t_{k+1})}
$$

$$
= \begin{bmatrix} 0, & -2\hat{b}^-(t_{k+1})\hat{b}^T_{meas}(t_{k+1}), & 0 \end{bmatrix}
$$

(31)

which is a 3x9 matrix.

The fixed-gain extended Kalman filter avoids the complexity of eq. 26-31 and the considerable computational burden of calculating the state transition matrix via numerical differentiation of $f[\ ]$. It can do this because fixed gains can stabilize the periodic observer associated with the periodic linearized system in eq. 16 and 17. Floquet analysis confirms this assertion. These gains have been calculated using a sub-optimal periodic observer theory that is similar to the sub-optimal control theory found in Anderson and Moore [9]. The algorithm for calculating such gains is very complicated and slow, but executes off line. Discussion of its theory is omitted, but results using this filter are presented below.

3.3 Filter Tuning

Filter tuning has two goals, timely convergence to an accurate estimate and maximum accuracy of the estimate. Filter tuning is possible through the selection of Q, R, and $P^-(t_0)$. Q and R determine the trade-off between the filtering of measurement noise and the rapid tracking of disturbance noise-induced state variations. $P^-(t_0)$ determines the rapidity of the initial filter convergence. In steady state, Q and R also

determine the filter stability as a by-product of the measurement noise/disturbance noise trade-off. $P^-(t_0)$ has no effect on the steady state performance of the filter.

The filter needs to have a rapid initial convergence because its mission is to accurately determine attitude with one orbit's worth of magnetometer data and poor initial attitude estimates. This means either a large $P^-(t_0)$ compared to R or a large Q compared to R. For the extended Kalman filter, which is inherently a time-varying filter, the former method has been used to achieve rapid initial convergence. The latter method has been used for the fixed-gain filter because it is a steady-state filter; it has no $P^-(t_0)$. This points to one advantage of time-varying filters: they allow rapid convergence without sacrifice of steady state filtering optimality.

Optimal steady-state tuning of the filter is important to achieving the accuracy goal of $\pm\ 1°$. The levels of measurement noise and disturbance torque that are present in a real S/C system make this a challenging goal. For the extended filter, this tuning is achieved by setting Q and R to magnitudes representative of the expected disturbance inputs and measurement noise. For the one case where detailed error analysis has been done, the disturbance torque level has been based on models of atmospheric drag torque and solar radiation pressure torque. The measurement noise has been based on magnetometer digitization error. The measurement error is actually much larger due to analog magnetometer noise and field model errors, so the filter is somewhat over sensitive to measurement noise in this case.

4 Evaluation of Filter Performance: Objectives and Tools

4.1 Filter Performance Criteria

There are two criteria for satisfactory filter performance: Does it converge? How accurate are its estimates of the S/C attitude? Because the system is nonlinear, filter stability is not guaranteed for large initial errors in the state estimate. Furthermore, the rate of convergence of stable filters is important because of the mission requirements; convergence must be achieved within one orbit. The importance of estimation accuracy is self evident. To evaluate the filter with respect to these two criteria is the objective of the test procedures that are outlined below. Results are reported in Section 5.

352

4.2 Analytical Tools for Performance Evaluation

The linearized model of the S/C's attitude dynamics, eq. 16 and 17, provides a valuable tool for analyzing filter stability and accuracy. The corresponding discrete-time model takes the form

$$x_{k+1} = \Phi(t_{k+1},t_k)x_k + \zeta_k \quad \text{for } k = 0,...,N\text{-}1 \quad (32)$$

$$y_k = H(t_k)x_k \quad \text{for } k = 1,...,N \quad (33)$$

where N is the number of magnetometer samples per orbit. Because the system in eq. 16 and 17 is periodic with the orbital period, the above discrete-time system is periodic with period N.

System observability from one orbit's worth of data can be analyzed by computation of the 1-orbit observability Gramian [10]. This Gramian is

$$\vartheta = \sum_{k=1}^{N} \Phi^T(t_k,t_0)H^T(t_k)H(t_k)\Phi(t_k,t_0) \quad (34)$$

If this 9x9 matrix is nonsingular, then the system is observable from one orbit's worth of data and there is hope for constructing a filter that converges in one orbit. One would expect this matrix to approach singularity with decreasing orbital inclination. A study of this dependence would map out the inclinations where the filters under consideration can be applied. The only observability Gramians computed for this study, however, were for the two orbital inclinations mentioned in Section 2, 43° and 57°.

Filter stability can be analyzed by applying Floquet theory to the discrete-time model of the steady-state filter. The steady-state, linear-model filter gains, K_k for k = 1...N, are periodic with period N; so, the filter itself is periodic. The one-orbit state transition matrix of the filter becomes:

$$\Phi_{CL}(t_N,t_0) = [I - K_N H(t_N)]\Phi(t_N,t_{N\text{-}1})...[I - K_1 H(t_1)]\Phi(t_1,t_0)$$

$$= \prod_{k=0}^{N\text{-}1}[I - K_{k+1}H(t_{k+1})]\Phi(t_{k+1},t_k) \quad (35)$$

where the CL subscript means closed-loop in the sense that the open-loop system state transition matrix and the filter gains are factored into this expression to give the 1-orbit observer error state transition matrix. The eigenvalues of $\Phi_{CL}(t_N,t_0)$ must all have magnitudes less than unity for filter stability, and the smallness of the eigenvalue magnitudes indicates the rate of convergence.

Accuracy of the filter can also be studied with the linearized, steady-state filter model. Given measurement and disturbance noise covariance matrices, R and Q, the periodic linear-filter covariance, $P^+(t_k)$ for k = 0...N, can be determined from the following linear system of equations:

$$P^+(t_{k+1}) = [I - K_{k+1}H(t_{k+1})][\Phi(t_{k+1},t_k)P^+(t_k)\Phi^T(t_{k+1},t_k)$$

$$+ Q][I - K_{k+1}H(t_{k+1})]^T$$

$$+ K_{k+1}R K^T_{k+1} \quad \text{for } k = 0,...,N\text{-}1$$

$$P^+(t_N) = P^+(t_0) \quad (36)$$

The Q and R matrices used here must be the best estimates of the actual disturbance and noise covariances, whereas those used in determining the K_{k+1} in an optimal or sub-optimal filter calculation may differ from the best estimates for various reasons. The 1-orbit average of the covariance yields the mean square filter accuracy:

$$P_{MS} = \frac{1}{N}\sum_{k=0}^{N\text{-}1} P^+(t_k) \quad (37)$$

which is a good measure of the effects of random disturbances and measurement noise on the filter accuracy.

4.3 Simulation Testing

Simulation testing is an important complement to analysis for purposes of filter evaluation. Nonlinearities may cause the filter to diverge for large initial attitude errors. Systematic errors such as parameter uncertainty or biases may degrade stability or accuracy or both. Linear analysis cannot evaluate these effects, but simulation can.

Each simulation test has two parts, a simulation and a filter. The simulation starts with an "actual" initial state, $x(t_0)$, and integrates the S/C attitude dynamics equations, eq. 1-3 or eq. 16, to produce a simulated "actual" state time history, $x(t_k)$ for k = 0,1,2,... It also simulates the magnetometer measurements to produce a measurement time history, $b_{meas}(t_k)$ for k = 0,1,2,... The filter takes these simulated magnetometer measurements, combined with initial estimates of the state and covariance, and produces an estimate of the state time history, $\tilde{x}^+(t_k)$ for k = 0,1,2,... Evaluation of the filter is accomplished by comparing the
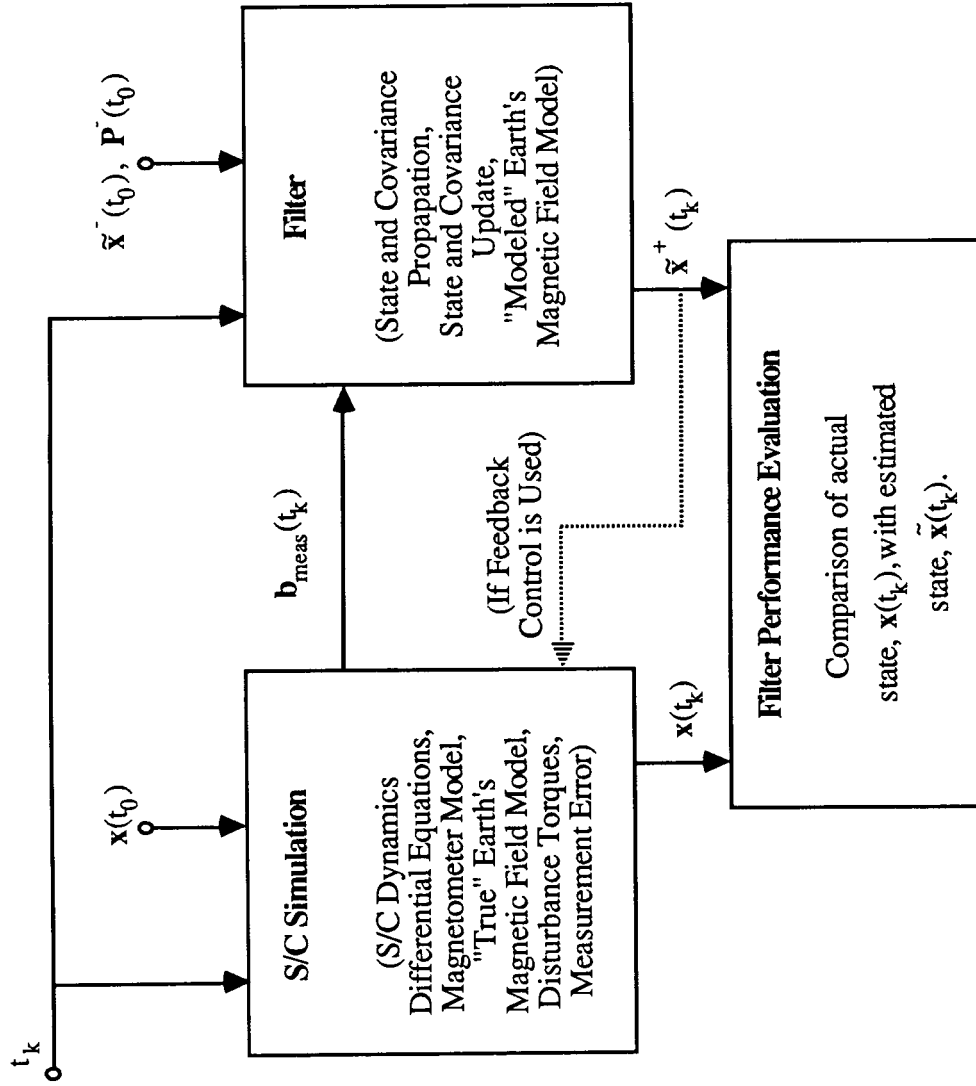
Figure 2. Block Diagram of Simulation Test Structure and Information Flow.

The block diagram contains the following labeled elements:

S/C Simulation

(S/C Dynamics
Differential Equations,
Magnetometer Model,
"True" Earth's
Magnetic Field Model,
Disturbance Torques,
Measurement Error)

Filter

(State and Covariance
Propapation,
State and Covariance
Update,
"Modeled" Earth's
Magnetic Field Model)

Filter Performance Evaluation

Comparison of actual
state, $x(t_k)$, with estimated
state, $\tilde{x}(t_k)$.

$x(t_0)$

$\tilde{x}^-(t_0)$, $P^-(t_0)$

t_k

$b_{meas}(t_k)$

(If Feedback
Control is Used)

$\tilde{x}^+(t_k)$

$x(t_k)$

"actual" state time history with the estimated state time history. Figure 2 depicts this two-part process, the flow of information between the parts, and the information used in evaluating the filter. For a good filter $\tilde{x}^+(t_k)$ will converge quickly to $x(t_k)$ and stay near it despite large discrepancies between $\tilde{x}^-(t_0)$ and $x(t_0)$ and despite disturbance torques, measurement noise, and modeling error.

Several properties of the filter have been evaluated with this simulation test scheme: the ability of the filter to converge from large initial attitude or rate errors, the ability of the filter to estimate constant disturbance torques, filter accuracy in the face of random disturbance torques and measurement noise, filter accuracy in the face of parameter errors in the S/C attitude dynamics model, and filter accuracy in the face of magnetometer biases. The attitude, rate, and torque estimation capabilities have been tested simply by running the nonlinear simulation and the filter starting each with different initial conditions (remember, disturbance torque is treated as a state). Filter accuracy in the face of random inputs has been evaluated by linear analysis, and for verification purposes, by simulation. The latter has been done by using time-varying measurement noise and disturbance torque models in appropriate parts of the simulation.

Attitude-dynamics parameter errors and measurement biases are systematic errors. Evaluation of their effects is tricky. They may or may not affect filter stability. They will certainly affect accuracy. The method used to evaluate these effects has been to simulate with one attitude dynamics/measurement process model and filter with a different model, the difference being the particular systematic error under consideration. The filter's stability and accuracy are then evaluated by comparison of the simulated "actual" state time history with the estimated state time history. Convergence is evaluated by comparing these two time histories for the first orbit. Accuracy is evaluated by taking the root mean square value of the difference between the estimated state and the "actual" state for all subsequent orbits. Convergence and accuracy are both dependent on the magnitude of the modeling errors. They may also be dependent on the magnitude of the S/C librations. Therefore, correct sizing of the errors and of the S/C libration amplitude is critical to correct analysis of these effects.

5 Filter Performance Results

5.1 Convergence

Filter convergence can be achieved if the 9x9 1-orbit observability Gramian (eq. 34) is nonsingular. This has been found to be the case for low Earth orbits at both 43° and 57° inclination. For S/C 1 of Table 1 in the 43° orbit, the 50-samples-per-orbit Gramian has a ratio of minimum eigenvalue to maximum eigenvalue of 8×10^{-7}. If the constant disturbance torque columns and rows are omitted, then the ratio increases to 3×10^{-6} for the resulting 6x6 sub-Gramian. For S/C 2 operating in the 57° orbit, a full Gramian has a minimum to maximum eigenvalue ratio of 6×10^{-13} (still nonzero in double precision arithmetic); whereas, the sub-Gramian for observing just the angles and the rates has an eigenvalue ratio of 6×10^{-8}. Thus, the filter can be made to converge in one orbit. The torques are less observable than the angles and rates. All cases are observable, but the second S/C case is less observable than the first, probably due more to the difference in S/C dynamic properties than to the difference in orbit.

The magnitudes of the eigenvalues of the one-orbit filter state transition matrix (eq. 35) are direct measures of stability. These have been computed for the time-varying filters in steady-state and for the fixed-gain filters. The fixed gain filter case that has been analyzed in detail involves S/C 1 operating in the 43° orbit at 50 magnetometer samples per orbit. The lowest achievable maximum eigenvalue magnitude for the fixed-gain 1-orbit state transition matrix has been 0.32, which indicates adequate stability but slow convergence. The time-varying extended Kalman filters do better, even in steady state. Their 1-orbit state transition matrices have maximum eigenvalue magnitudes typically less than 0.20. The initial convergence of these time-varying filters is even better than this steady-state result indicates because of the high values selected for the initial covariance matrix, $P^-(t_0)$.

This rapid initial convergence is indicated clearly in Fig. 3 and 4 for S/C 2 in the 57° orbit. The time-varying filter used in this case operates on about 300 magnetometer samples per orbit; all of the S/C-1 cases used 50 samples per orbit. In these figures as in most of the remaining figures, the "actual" value, the estimated value, and the estimation error for a particular quantity are all plotted together on a single graph. The orbital period is a little over 5,000 sec, so
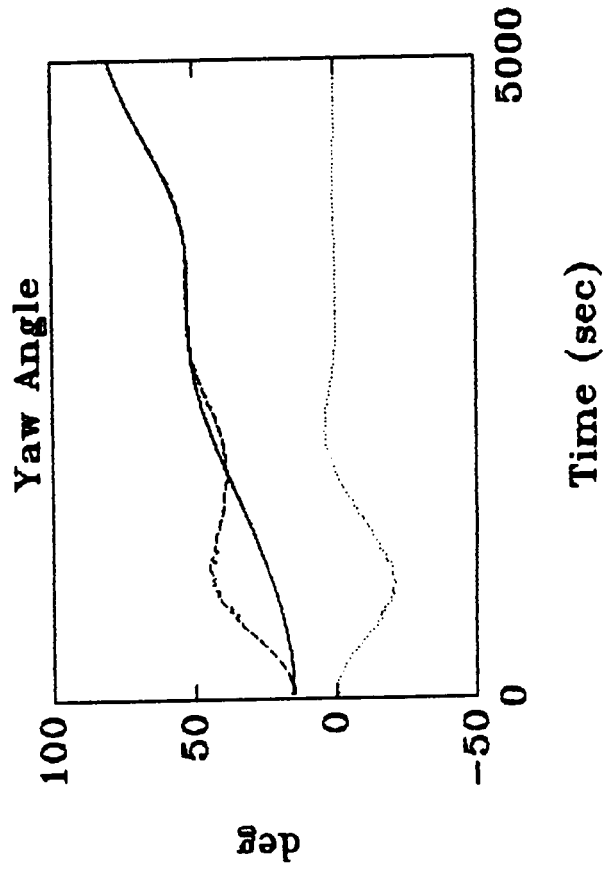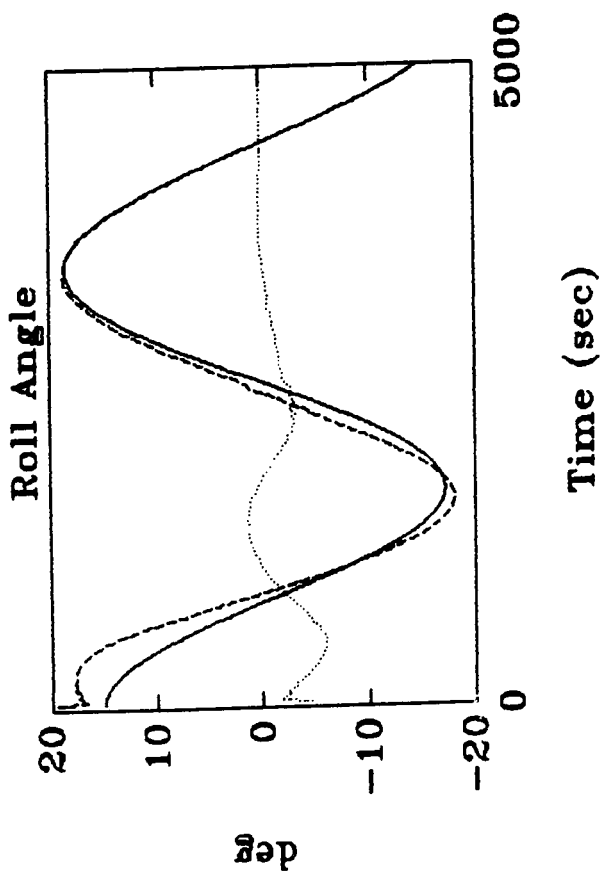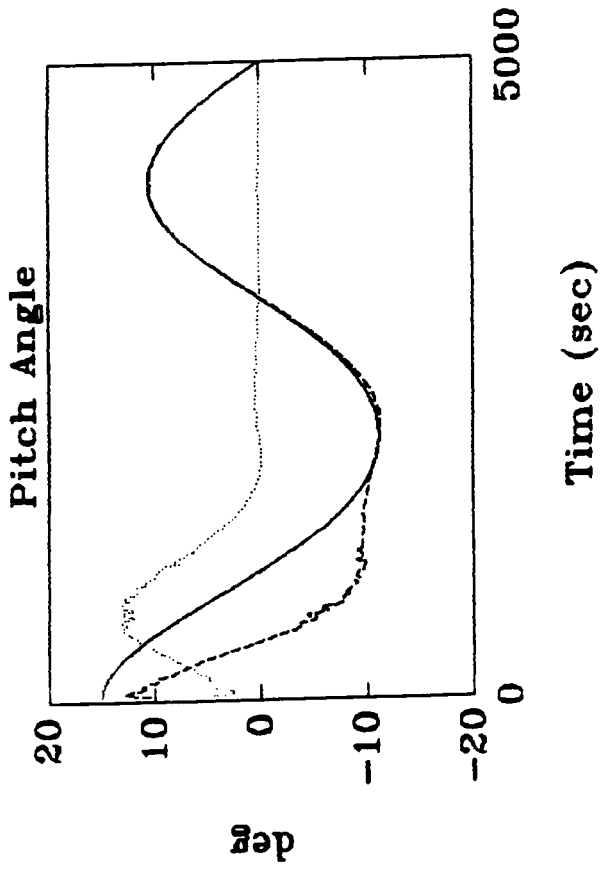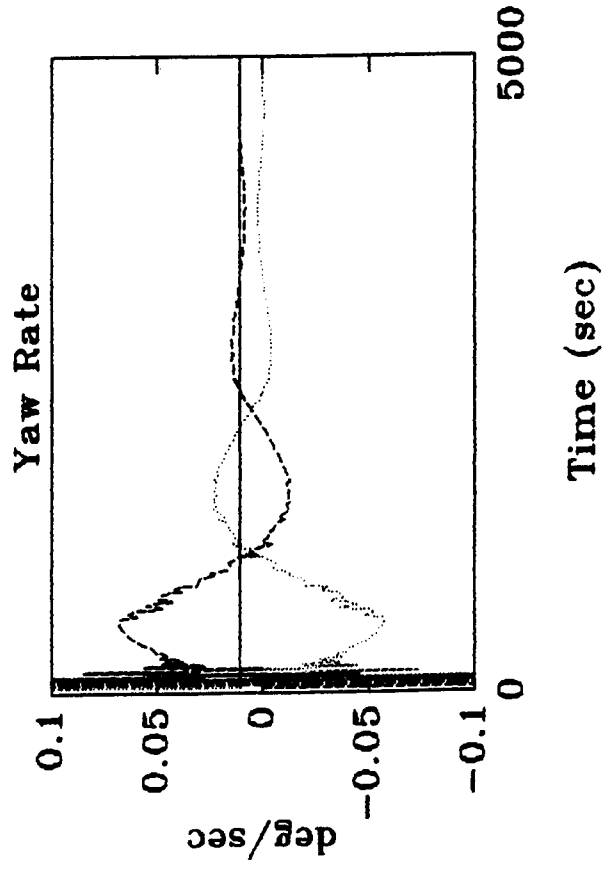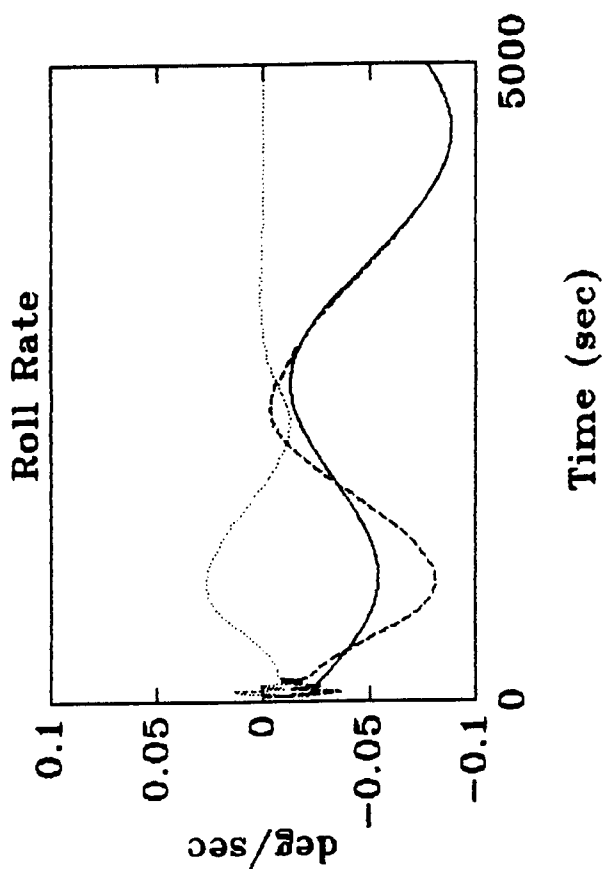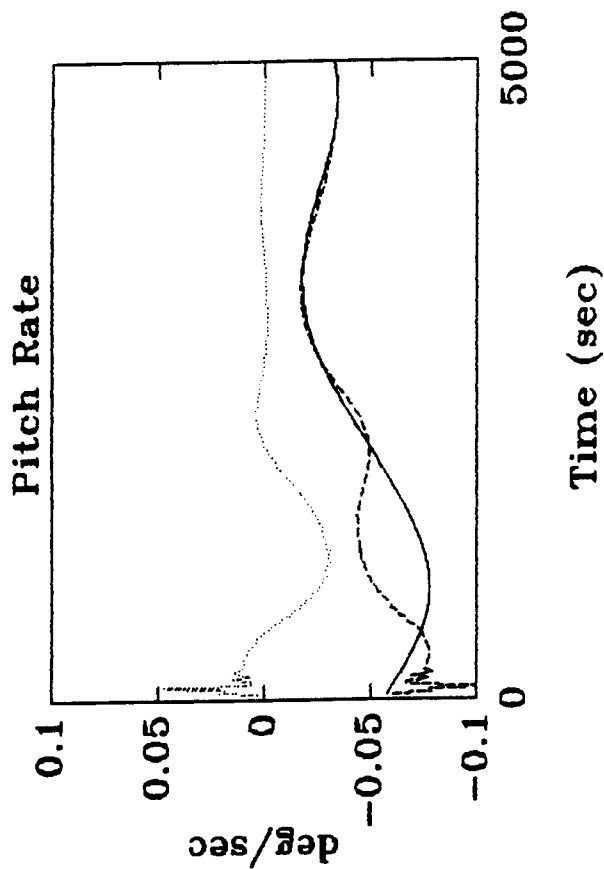
355

Figure 3. Moderate Initial Attitude Estimation Error,
15° on All Three Axes.
(Attitude Time Histories for Spacecraft 2)

Solid   = Actual

Dashed = Estimate

Dotted = Error

Figure 4. Moderate Initial Attitude Estimation Error, 15° on All Three Axes. (Attitude Rate Time Histories for Spacecraft 2)

Solid = Actual

Dashed = Estimate

Dotted = Error

357

convergence takes place within about half an orbit. The initial error is 15° each in roll, pitch, and yaw with no initial rate errors. The initial filter jitter in the rate estimates (Fig. 4) arises because of the high initial filter gains. Perhaps these are too high. Essentially, the filter is trying to take first and second derivatives to get rate and torque information.

Figures 5 and 6 further demonstrate the ability of the time-varying extended Kalman filter to converge. In the case of Fig. 5, a 0.017 N-m aerodynamic pitch torque is acting on S/C 1, which is in the 43° orbit. The filter converges to a correct estimate of this disturbance torque in just about one orbit (5,700 sec). Figure 6 corresponds to the same S/C-orbit case with large initial errors in the attitude estimate; the total initial rotational error is 45°. The filter successfully converges in about one orbit, despite the increased significance of nonlinearities.

Figure 7 also depicts estimation of the attitude of S/C 1 in the 43° orbit, but the estimates have been generated by a fixed-gain filter. These estimates converge from moderate initial errors (11° in all three axes), but convergence is slow. After almost two orbits errors on the order of 1° still persist. This makes sense in light of the large maximum filter eigenvalue. This convergence rate is too slow for the filter's intended purpose. Time-varying filters are preferable to fixed gain filters because of the ability to achieve faster initial convergence by increasing $P^-(t_0)$.

5.2 Steady-State Error Analysis

Error analysis has been done for S/C 1 operating in the 43° orbit. This error analysis combines the linear analysis technique described in Section 4.2 with the simulation technique described in section 4.3. The linear technique has been used for random errors, the simulation technique for systematic errors. The final error budget combines the two in a square-root-of-the-sum-of-the-squares (RSS) sense.

The random effects considered are time varying solar and atmospheric drag disturbance torques (constant disturbance torques do not affect the error because they are estimated as part of the filter state vector), random magnetometer measurement error and digitization error, and random or high-order International Geomagnetic Reference Field (IGRF) model error. The disturbance covariance matrix magnitude is based on the results of a solar torque and aerodynamic torque analysis for the S/C. The

magnetometer random error is based on a 5 mGauss spec for its accuracy and a 12-bit digitization. The field error was set at 0.41° rms per axis based on experience with the IGRF model data [6].

Figure 8 depicts a simulation of the effects of one of the random errors, magnetometer digitization error. Initial convergence is hardly affected by this random process; convergence to within 0.5° still occurs within one orbit. Afterwards, random effects dominate the error signal. This calls for proper selection of the Q and R filter matrices for optimal steady state filter performance.

The systematic error magnitudes have been derived from typical S/C 1 specifications. Errors of 2% in the magnitudes of the principal moments of inertia were used. This number is based on the possible variability of the boom lengths and tip mass weights. Errors of 1.5° for the principal axis orientations were used to generate the cross product of inertia errors. This magnitude is based on angular accuracy specifications for the booms. A 5% error for the pitch wheel angular momentum was used, based on hardware specifications. A 1.4 N-m-sec error for the magnetically-anchored damping constant was assumed, a typical level of variability on orbit. A 3 mGauss bias error per axis for the magnetometer was used based on a typical magnetometer spec.

Representative libration magnitudes were used for the error budget simulations: 0.4° peak-to-peak roll angle oscillations, 4.0° peak-to-peak pitch angle oscillations, and 1.6° peak-to-peak yaw angle oscillations. These libration magnitudes are based on analysis of typical S/C 1 motions on orbit.

Figures 9 and 10 are typical of the simulation/filter time histories that have been used to evaluate the effects of systematic errors. The simulation in Fig.9 corresponds to discrepancies between the filter model and the simulation model of 1.5° in all three principal axes. The discrepancy corresponding to Fig. 10 is 1.4 N-m-sec in the passive damping constant. The filter converges in both these cases. As per the analysis description at the end of Section 4.3, the steady state error contributions have been taken to be the post-one-orbit rms error values.

Table 2 summarizes the error budget for this case. According to this analysis the filter meets the 1° attitude
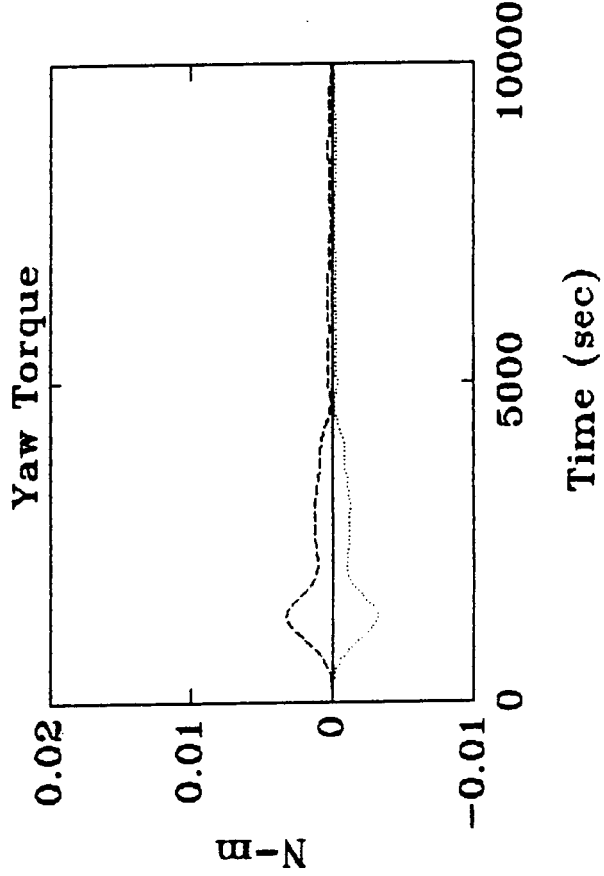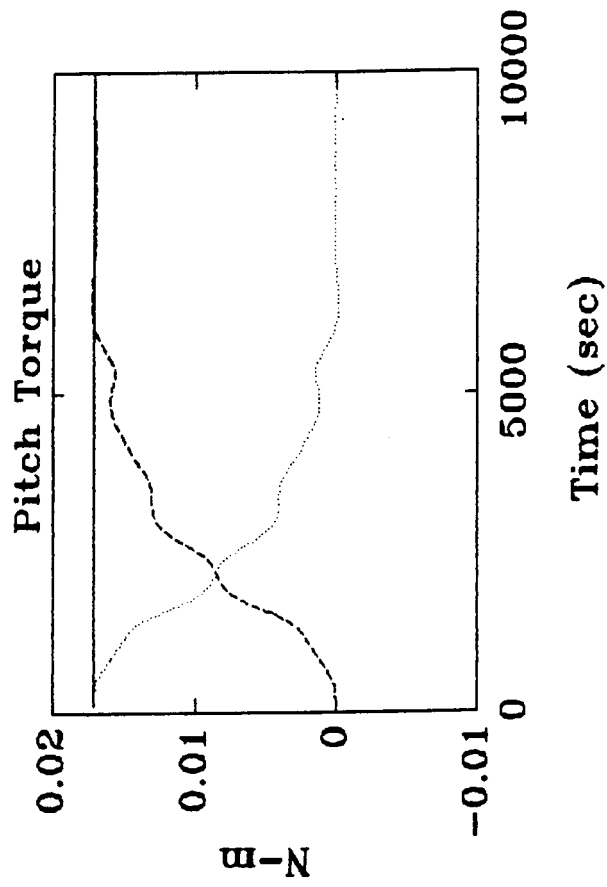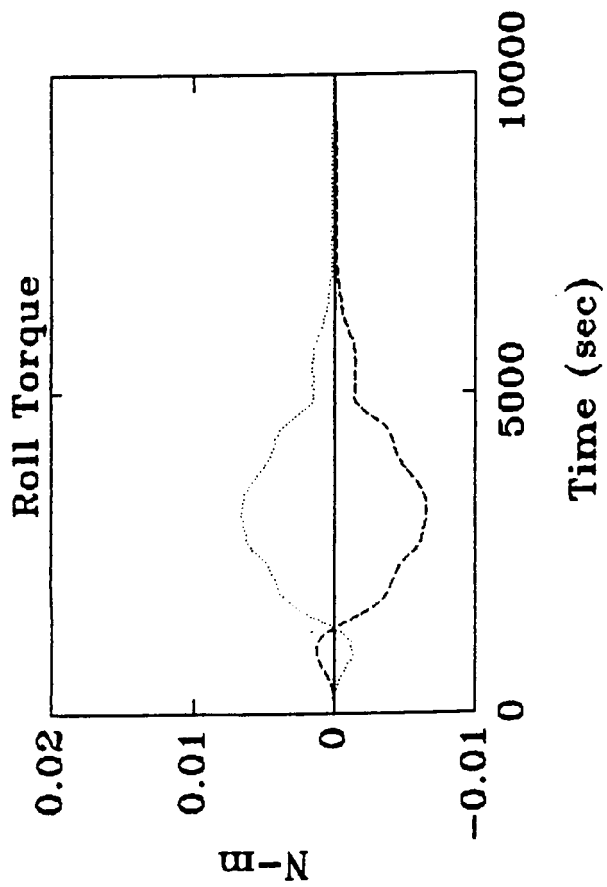
Figure 5. Initial Disturbance Torque Estimation Error, 0.017 N-m on the Pitch Axis. (Disturbance Torque Time Histories for Spacecraft 1)
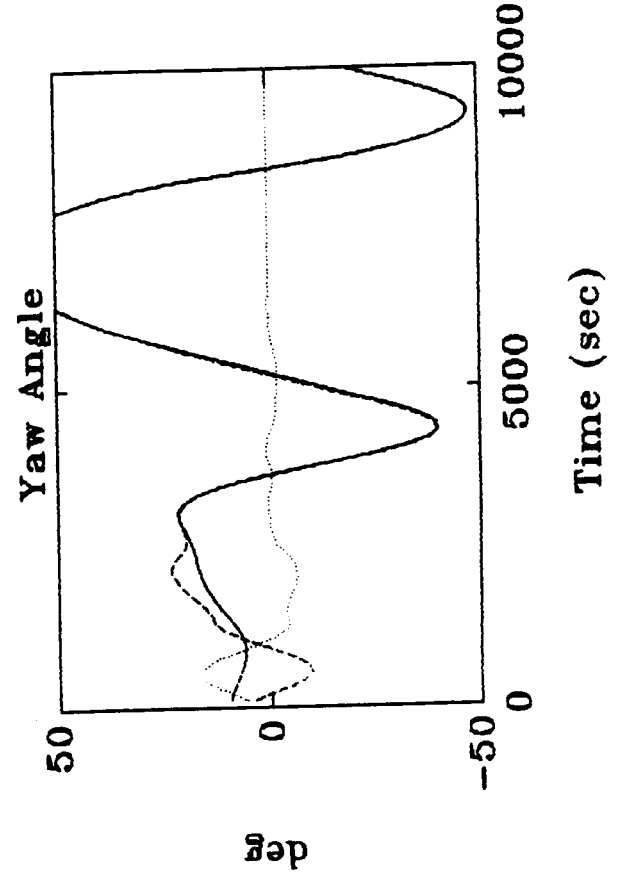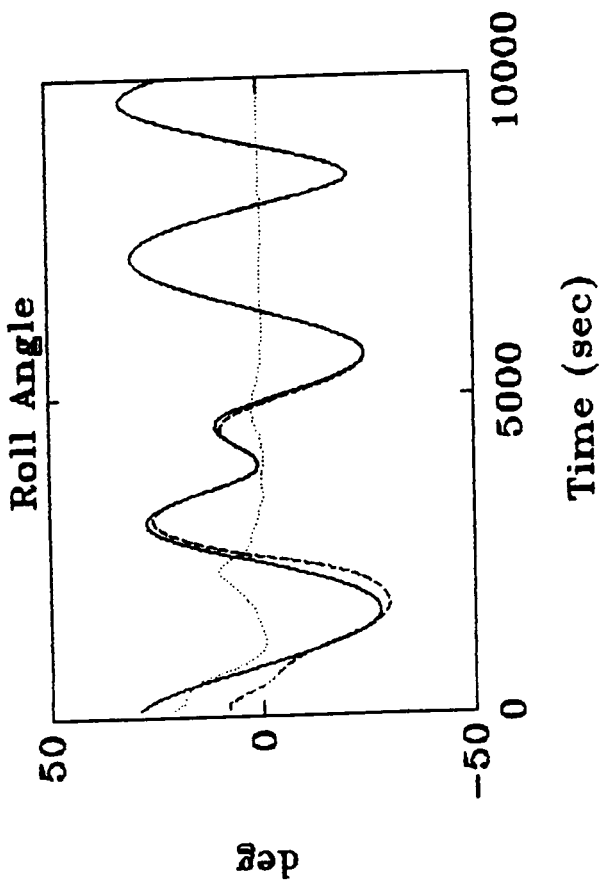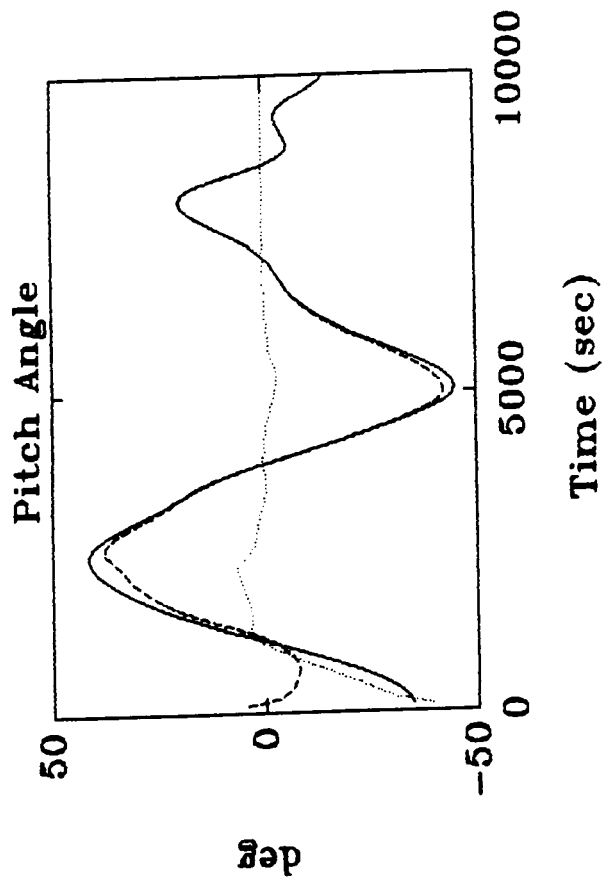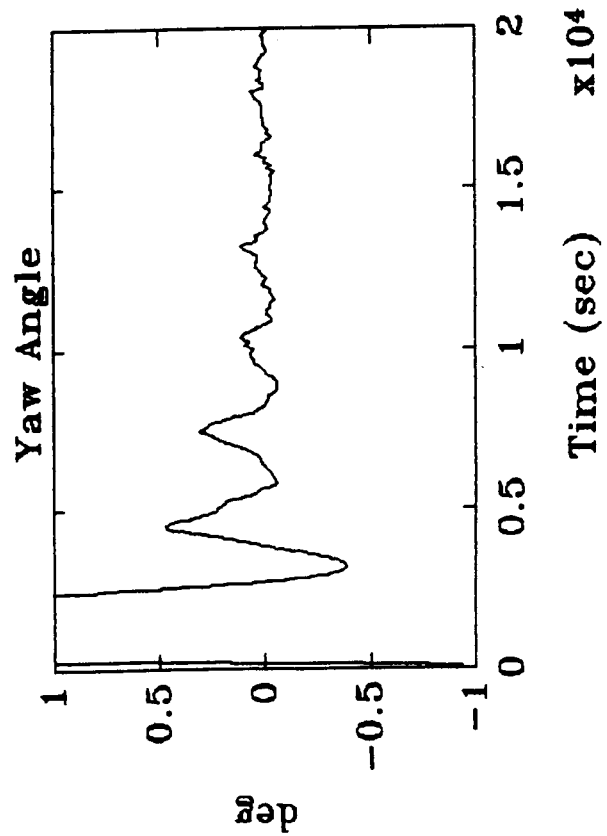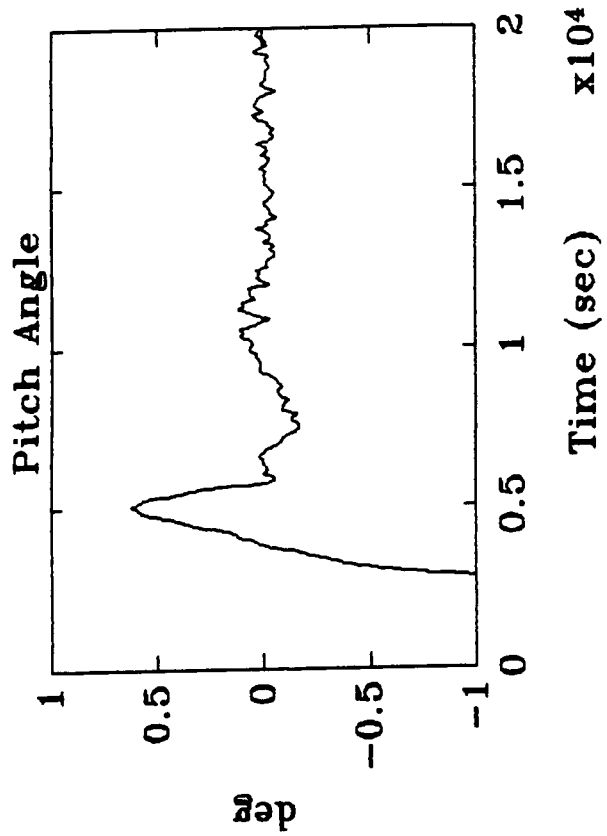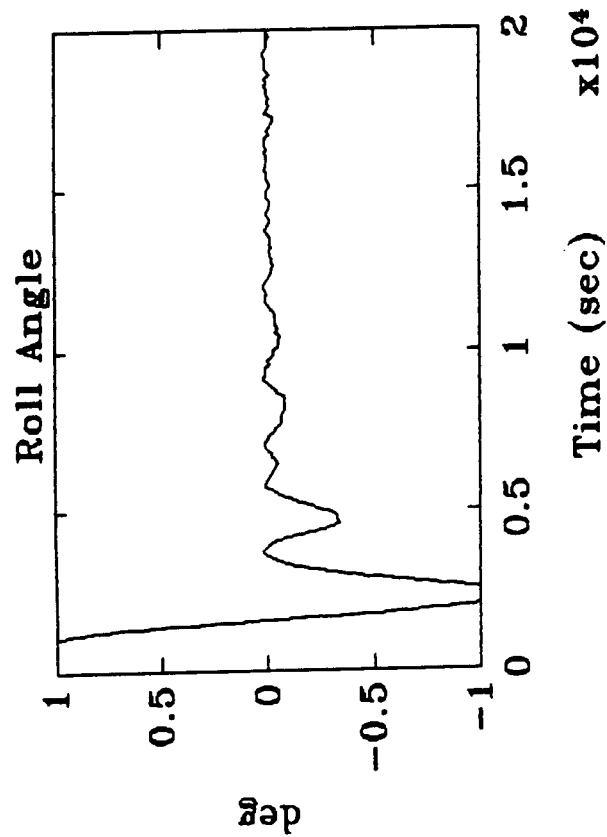
Solid = Actual

Dashed = Estimate

Dotted = Error

Figure 6. Large Initial Attitude Estimation Error, 45° Total Euler Rotation Angle. (Attitude Time Histories for Spacecraft 1)
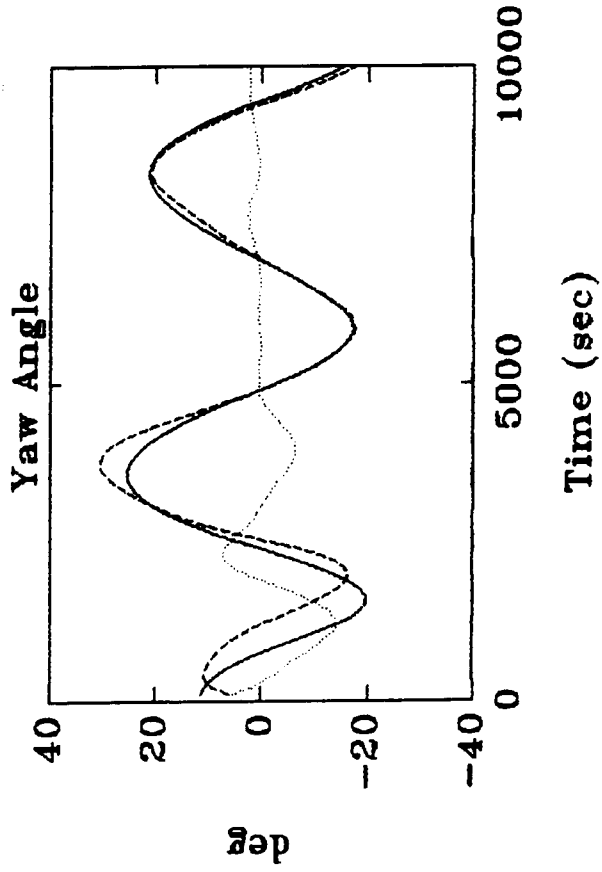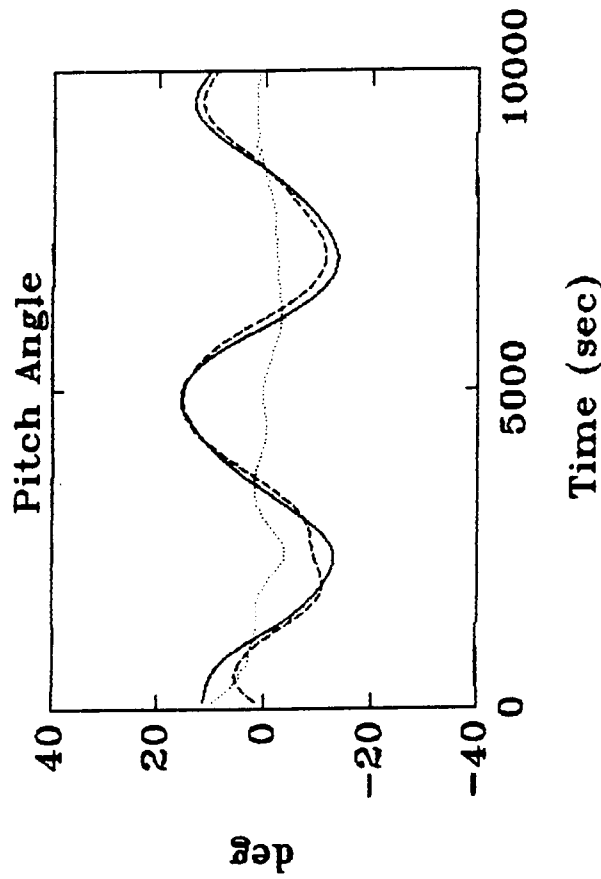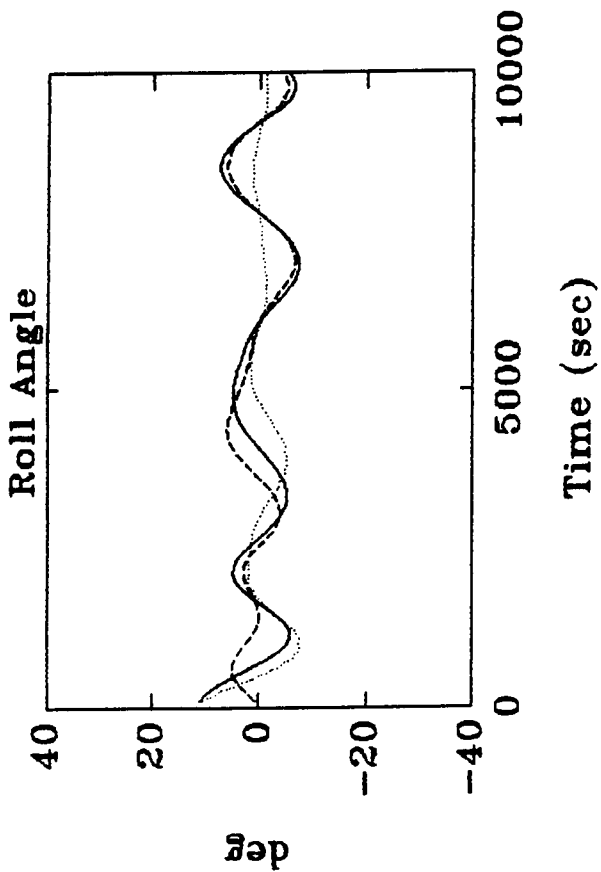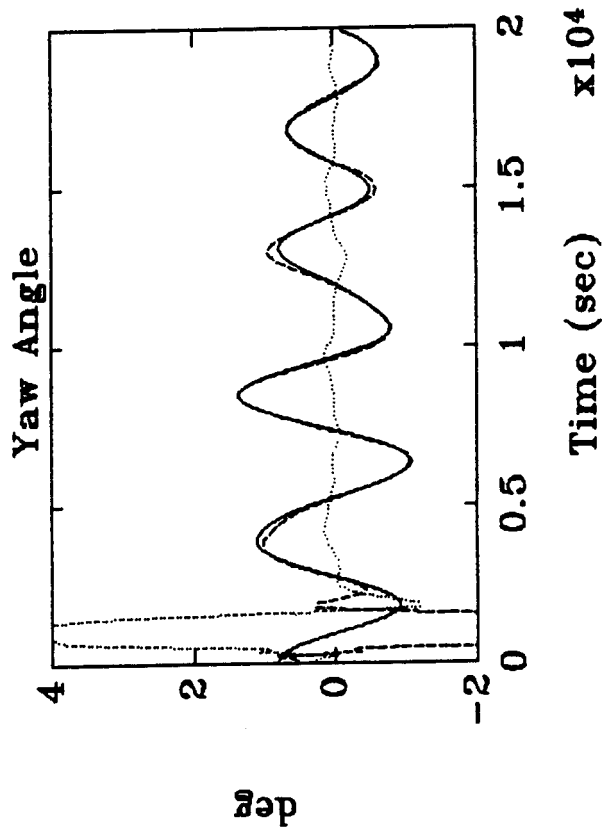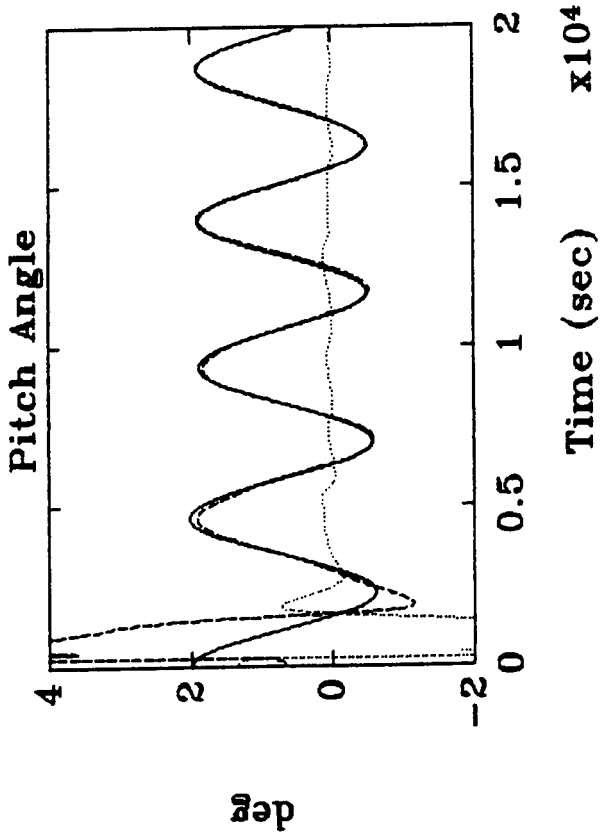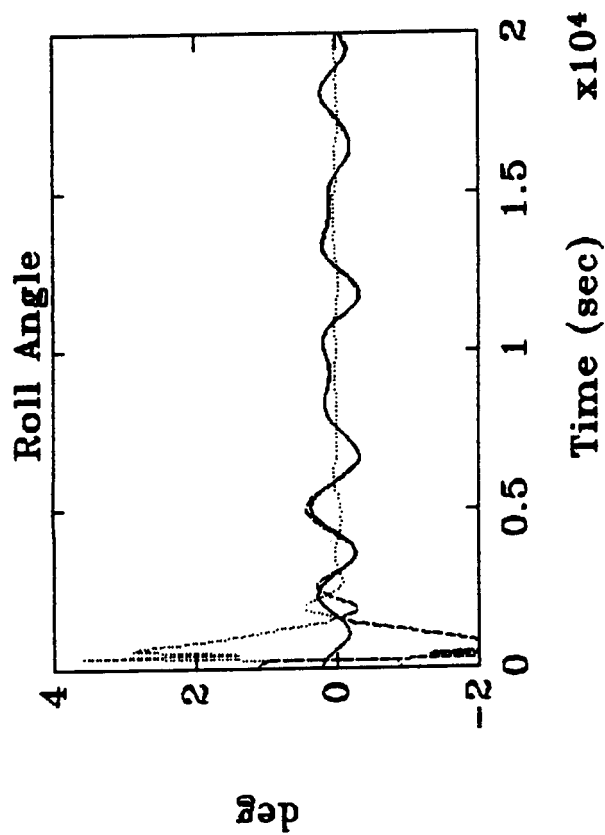
Figure 7. Fixed-Gain Filtering with a Moderate Initial Attitude Estimation Error, 11° on All Three Axes. (Attitude Time Histories for Spacecraft 1)

Solid   = Actual

Dashed = Estimate

Dotted = Error

Figure 8. The Effect of Magnetometer Digitization Error
on Estimation Error.
(Attitude Error Time Histories for Spacecraft 1)

Figure 9. The Effect of Product-of-Inertia Modeling Error on Estimation Error. (Attitude Time Histories for Spacecraft 1)
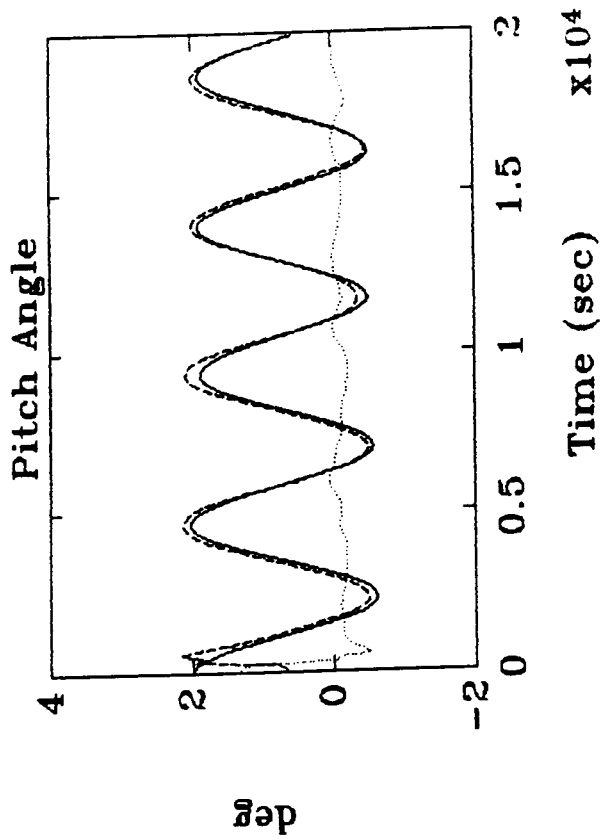
Solid = Actual
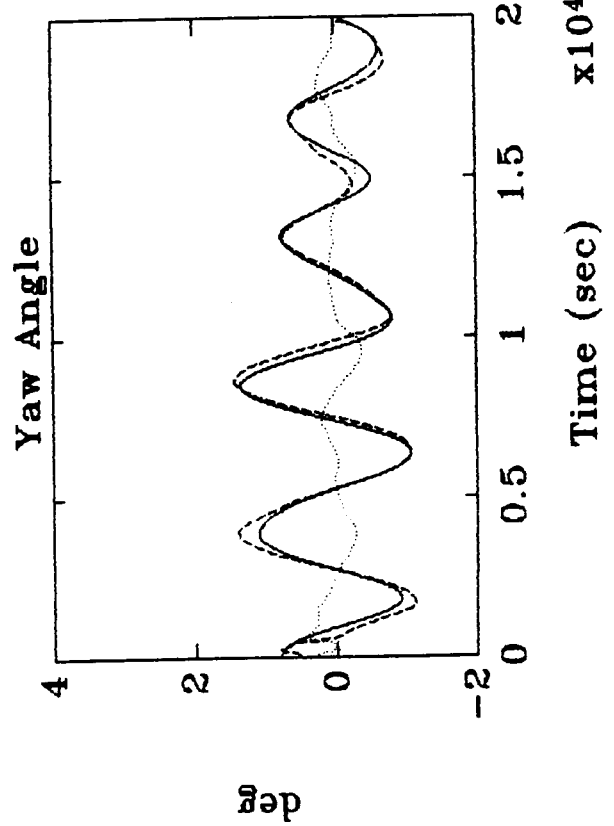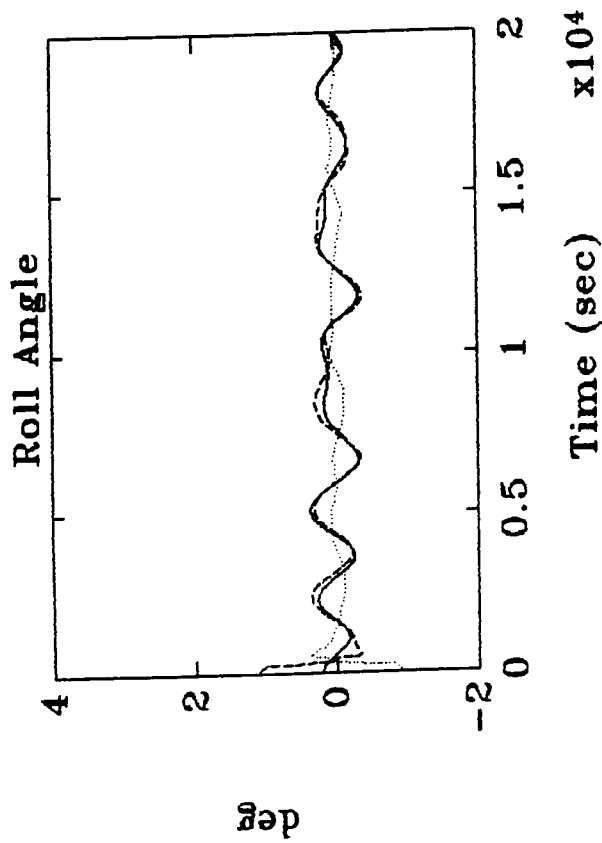
Dashed = Estimate

Dotted = Error

Figure 10. The Effect of Damping-Constant Modeling Error on Estimation Error. (Attitude Time Histories for Spacecraft 1)

Table 2

Error Analysis Results for S/C 1

Error Source	RMS Value (deg)		
	Roll	Pitch	Yaw
2% Roll-Axis Moment of Inertia Error	0.015	0.021	0.026
2% Pitch-Axis Moment of Inertia Error	0.016	0.027	0.033
2% Yaw-Axis Moment of Inertia Error	0.006	0.011	0.013
1.5° Principle Axes Skew Error (in all three axes)	0.017	0.041	0.063
1.4 N-m-sec Magnetically-Anchored Damping Error	0.056	0.112	0.159
5% Pitch Wheel Momentum Error	0.008	0.014	0.017
3 mGauss Magnetometer Bias Error	0.231	0.369	0.568
Random Measurement and Field Model Error	0.345	0.789	0.816
Random Disturbance Torque Error	0.009	0.016	0.023
RSS Error	0.420	0.880	1.010

knowledge goals in roll and pitch, but misses slightly in yaw. These are 1-σ numbers.

The main contributors to the errors are uncertainty in the IGRF model and magnetometer accuracy limitations. Increasing the magnetometer accuracy by a factor of 5 would decrease the RSS error by 40 %. Increasing the field model accuracy would also have a significant beneficial impact on the error budget. Decreasing the error in the knowledge of the magnetically-anchored damping constant would further improve the filter accuracy. Increasing the **R** weighting in the filter might also improve things. The current low **R** value causes the filter to rely too heavily on each magnetometer measurement, hence the large effects on accuracy of measurement-type errors and the small effects of disturbance torque errors and modeling errors.

A significant result of this analysis is the relative insensitivity of the filter to angular shifts in the S/C principal axes. The filter identifies these as disturbance torques and continues to achieve attitude accuracy on the order of 0.06° or better (neglecting other contributions to error), despite the 1.5° bias in the spacecraft attitude from that predicted by gravity gradient analysis and the modelled products of inertia.

5.3 Notes on Filter Performance

This filter runs relatively fast. It is able to perform one orbit's (50 samples) worth of filtering in about 3 minutes when operating on an INTEL 8088/8087-based personal computer with an 8 MHz clock rate. This time does not include the time to compute the IGRF model from a spherical expansion. In this particular filter implementation, the field calculation is done offline and the filter gets the resulting data from a table look-up.

Several problems having to do with convergence occur with this filter. One problem occurs in covariance initialization. For the case of "large" diagonal $P^-(t_0)$ and "small" **R**, if the initial magnetometer measurement occurs at t_0, then $P^-(t_0)$ is immediately updated to yield $P^+(t_0)$ before the first covariance propagation, and the filter converges. If the first magnetometer measurement occurs at t_1, on the other hand, then propagation to $P^-(t_1)$ occurs first, and the filter sometimes fails to converge. This is because $\Phi(t_1,t_0)$ has some large elements for long sampling periods, which makes $P^-(t_1)$ very high and results in very high filter gains. These cause divergence due to nonlinearities.

Another convergence problem occurs because of the quaternion update scheme in eq. 24 and 25. The argument of the square root in eq. 24 becomes negative for very large

initial attitude errors. This occurs because the linear update in eq. 24 does not recognize when it is asking for more angular correction than makes physical sense. This problem occurred for an initial angular error of 90° (the filter did converge for a 60° initial error). It also can occur when filter gains are too high -- $\mathbf{P}^-(t_0)$ or $\mathbf{Q}$ or both are too large relative to $\mathbf{R}$. A way to avoid this would be to add a further nonlinearity to the quaternion update to scale down updates that are physically unrealizable.

6 Conclusions

The modified extended Kalman filter described and analyzed in this paper can estimate 3-axis S/C attitude, attitude rate, and constant disturbance torques solely from 3-axis magnetometer measurements distributed over one orbit. The filter works for gravity-gradient stabilized S/C operating in inclined, low-Earth orbits. The filter can converge from initial attitude errors as large as 60° and can achieve a 1-σ attitude accuracy of 1° or better on all three axes.

Filter performance has been evaluated in two ways: by linear analysis of small perturbations from the nominal gravity-gradient orientation and by filtering of magnetometer data generated by a nonlinear S/C simulation. The linear analysis has confirmed the observability of the system and the stability of the filter. Also, it has predicted the inaccuracy induced by random disturbances and measurement noise. The nonlinear simulation has demonstrated the filter's ability to converge from large initial errors and has predicted the contributions of systematic errors to inaccuracy.

For one of the S/C-orbit cases considered, the most significant contributors to filter inaccuracy are magnetometer inaccuracy and inaccuracy of the knowledge of the Earth's magnetic field. Reduction of the filter's 3-σ attitude uncertainty to 1° could be achieved by use of a magnetometer with 1 mGauss 1-σ accuracy in combination with a model of the Earth's magnetic field that is accurate to 0.1° 1-σ. The necessary accuracy increases are about 3 times for the magnetometer and about 4 times for the field model. More accurate predictions of the passive magnetically-anchored damping factor would also improve filter accuracy.

7 Recommendations and Planned Follow-Up Work

Comparison of these accuracy and convergence results with flight test results is planned. There are plans to launch a satellite using this filter in the ground station as a back-up attitude determination system. The S/C will also carry optical attitude determination instruments, which are more accurate than the magnetometer. Comparison of filter attitude with the attitude determined by the more accurate system will provide a bench mark for its evaluation. A post-launch action that will be considered for improvement of the filter accuracy is on-orbit magnetometer calibration and bias determination similar to what was done in Ref. 4.

A related application for magnetometer data filtering could be made on the autonomous satellite navigation problem. The observability of the attitude/trajectory system should be checked, this time retaining length information for the field vector. The system may be observable. In that case, a navigation system based solely on magnetometer measurements or on a combination of horizon sensor and magnetometer measurements would be theoretically possible.

The filter described in this paper coupled with a magnetometer and sufficient computer capacity can be used when a low-cost, light-weight, low-accuracy system for attitude determination is required.

Acknowledgements

References

1. 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, Sept.-Oct. 1982, pp. 417-429.

2. Potter, J.E., and Vander Velde, W.C., "Optimum Mixing of Gyroscope and Star Tracker Data, " **Journal of Spacecraft and Rockets**, Vol. 5, May 1968, pp. 536-540.

3. Gai, E., Daly, K., Harrison, J., and Lemos, L., "Star-Sensor Based Satellite Attitude/Attitude Rate Estimator," **Journal of Guidance, Control, and Dynamics**, Vol. 8, No. 5, Sept.-Oct. 1985, pp. 560-565.

4. Thompson, R.H., Neal, G.F., and Shuster, M.D., "Magnetometer Bias Determination and Spin Axis Attitude

Estimation for the AMPTE Mission," **Journal of Guidance, Control, and Dynamics**, Vol. 7, No. 4, July-Aug. 1984, pp. 505-507.

5. Heyler, G.A., "Attitude Determination by Enhanced Kalman Filtering Using Euler Parameter Dynamics and Rotational Update Equations," AIAA Paper No. A81-45832, AAS/AIAA Astrodynamics Specialist Conference, Lake Tahoe, Nevada, Aug. 3-5, 1981.

6. Wertz, J.R. ed., **Spacecraft Attitude Determination and Control**, D. Reidel Pub. Co., (Boston, 1978).

7. Bar-Itzhack, I.Y., and Oshman, Y., "Attitude Determination from Vector Observations: Quaternion Estimation," **IEEE Trans. on Aerospace and Electronic Systems**, Vol. AES-21, No.1, Jan. 1985, pp. 128-135.

8. Stengel, R.F., **Stochastic Optimal Control**, J. Wiley and Sons, (New York, 1986).

9. Anderson, B.D.O., and Moore, J.B., **Linear Optimal Control**, Prentice-Hall, (Englewood Cliffs, N.J., 1971).

10. Kailath, T., **Linear Systems**, Prentice-Hall, (Englewood Cliffs, N.J., 1980).