A Comparison of Linear Attitude Estimators

Renato Zanetti*
The University of Texas at Austin, Austin, Texas 78712

Kyle J. DeMars†
Missouri University of Science and Technology, Rolla, Missouri, 65409

This work compares the performance of linear attitude estimators; in particular, the classic multiplicative extened Kalman filter and unscented Kalman filter performances are compared to a recently introduced novel spacecraft attitude estimation algorithm. The new algorithm utilizes unit vector measurements and is also based on the unscented Kalman filter (UKF). The UKF, like the extended Kalman filter, is an approximation of the linear minimum mean square error estimator and employs a linear update with an additive residual. The standard formulation of the residual is given by the difference between the measurement and its mean. The recently proposed algorithm, on the other hand, utilizes a multiplicative residual, which is more consistent with the nature of unit direction measurements. The recent algorithm consistently defines attitude errors utilizing the Gibbs vector parameterization and computes averages and deviations consistently with attitude composition rules.

I. Introduction

This work compares three different attitude-specific implementations of the extended Kalman filter (EKF) using the quaternion-of-rotation [1, 2]. Attitude quaternions have a unit-norm constraint that can be enforced via the multiplicative extended Kalman filter (MEKF) [3], the additive extended Kalman filter (AEKF) [4], projection techniques [5], constrained Kalman filtering [6], or the unscented quaternion estimator [7] that is based on the unscented Kalman filter (UKF) [8].

Both the EKF and the UKF are linear estimators as the measurement update is constructed as a linear function of the measurement. Linear estimators necessitate knowledge of the first- and second-order moments of distributions, a non-trivial task for nonlinear measurements. The EKF approximates these moments through linearization around the mean, while the UKF utilizes statistical linearization [9]. The UKF approximates first-order moments (means) and second-order moments (covariances) accurate to at least second order [10].

This paper compares three linear attitude estimators: the MEKF, the unscented quaternion estimator, and a recent algorithm proposed in Ref. [11] in the presence of unit vector measurements. Usage of unit vectors as attitude measurements is a widely adopted technique [4], and various measurement models have been proposed, such as the QUEST model [12], the wide-field-of-view model [13], and the multiplicative measurement model [14]. This paper builds on prior work by the authors in [11] and [15].

The MEKF is based on the assumption that small attitude errors are locally a vector space, and hence are treated as vectors. The unscented quaternion estimator by Crassidis and Markley [7] uses an additive measurement model and an additive residual for the measurement update and the algebraic average of three-dimensional attitude error for the time propagation. Ref. [11] employs a multiplicative measurement model and a multiplicative residual [16, 17] for the update phase and quaternion averaging [18], for propagation.

II. The Multiplicative Extended Kalman Filter

We work with a dynamical system of the form

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

where $\mathbf{x}_{k-1}$ is the state at time t_{k-1}, $\mathbf{x}$ is the state at time t_k, and $\mathbf{w}_{k-1}$ is the zero-mean and white process noise with covariance matrix Q_{k-1}. The mean evolves as

$$\hat{\mathbf{x}}^- = \mathrm{E}\left\{\mathbf{f}(\mathbf{x}_{k-1}, \mathbf{w}_{k-1})\right\}$$

*Assistant Professor, Department of Aerospace Engineering and Engineering Mechanics, AIAA Associate Fellow
†Assistant Professor, Department of Mechanical and Aerospace Engineering, AIAA Senior Member

The MEKF's attitude state is a quaternio;, however, the attitude estimation error is not parameterized as a quaternion but rather with a three-dimensional attitude representation. Errors are assumed to be small (which is consistent with the linearization assumption above), and we take advantage of the fact that attitude is locally a vector space, leading to

$$\bar{\mathbf{q}} = \bar{\mathbf{q}}(\mathbf{e}_{\boldsymbol{\theta}}) \otimes \bar{\boldsymbol{q}} \tag{1}$$

where $\bar{\mathbf{q}}(\mathbf{e}_{\boldsymbol{\theta}})$ is the quaternion parameterization of the three-dimensional attitude error $\mathbf{e}_{\boldsymbol{\theta}}$. The MEKF propagates the estimate $\hat{\mathbf{x}}$ trough linearization around the mean

$$\hat{\mathbf{x}}^- \simeq \mathrm{E}\left\{\mathbf{f}(\hat{\mathbf{x}}_{k-1}, \mathbf{0}) + F_{k-1}\mathbf{e}_{k-1} + G_{k-1}\mathbf{w}_{k-1}\right\} = \mathbf{f}(\hat{\mathbf{x}}_{k-1}, \mathbf{0}) \tag{2}$$

since the process noise and estimation error are assumed zero mean and where F_{k-1} and G_{k-1} are appropriate partials. The estimation error covariance matrix P^- is then propagated as

$$P^- = F_{k-1} P_{k-1} F_{k-1}^{\mathrm{T}} + G_{k-1} Q_{k-1} G_{k-1}^{\mathrm{T}} \tag{3}$$

When a measurement $\mathbf{y}$ becomes available, its mean is given by

$$\hat{\mathbf{y}}^- = \mathrm{E}\left\{\mathbf{h}(\mathbf{x}, \mathbf{v})\right\}$$

where $\mathbf{v}$ is the measurement noise and $\mathbf{h}$ is a nonlinear measurement function. A linear update that minimizes the mean square estimation error is given by

$$\begin{aligned} \hat{\mathbf{e}}^+ &= \hat{\mathbf{e}}^- + K(\mathbf{y} - \hat{\mathbf{y}}^-) \\ P^+ &= P^- - K P_{yy} K^{\mathrm{T}} \\ P_{yy} &= \mathrm{E}\left\{(\mathbf{y} - \hat{\mathbf{y}}^-)(\mathbf{y} - \hat{\mathbf{y}}^-)^{\mathrm{T}}\right\} \end{aligned}$$

where the estimate of the prior estimation error $\hat{\mathbf{e}}^-$ is zero by construction and K is the Kalman gain

$$\begin{aligned} K &= P_{xy} P_{yy}^{-1} \\ P_{xy} &= \mathrm{E}\left\{(\mathbf{y} - \hat{\mathbf{y}}^-)\mathbf{e}^{\mathrm{T}}\right\} \end{aligned}$$

The MEKF approximates $\hat{\mathbf{y}}$, P_{yy} and P_{xy} via linearization around the mean, i.e.

$$\mathbf{y} - \hat{\mathbf{y}}^- \simeq H\,\mathbf{e} + L\,\mathbf{v} \tag{4}$$

where H and L are appropriate partials.

In summary, the MEKF treats attitude errors and unit vector direction measurements as vectors (it adds and subtracts them). Notice that for unit vector measurements, contingent upon the linearization assumption, the MEKF additive residual approach is equivalent to the multiplicative residual approach [17], and the two are effectively identical to processing the bearing angles directly[19].

III. The Quaternion Unscented Kalman Filter

This section presents the algorithm by Crassidis and Markley [7]. The general unscented background material follows the presentation in [15].

A. The Unscented Transform

Starting from a nonlinear random vector transformation

$$\mathbf{z} = \mathbf{g}(\mathbf{x})$$

the unscented transform (UT) approximates the transformation of the mean $\mathbf{m}_z$ and covariance P_{zz} of the output, $\mathbf{z}$ with a deterministic set of points with discrete probability, the so-called sigma-points. The sigma-points are chosen such that $\mathbf{m}_x$ and P_{xx} are maintained exactly. The sigma-points are then applied as inputs to the nonlinear function to yield nonlinearly transformed sigma-points, which can then be used to approximate a nonlinear transformation of the output mean and covariance, $\mathbf{m}_z$ and P_{zz}.

The K sigma-points are denoted by $\mathcal{X}_i$ and their associated weights by w_i, where $i \in \{1, \ldots, K\}$ and $\sum_{i=1}^{K} w_i = 1$. We then apply the transformation on each of the sigma points to obtain

$$\mathcal{Z}_i = \mathbf{g}(\mathcal{X}_i) \quad \forall \;\; i \in \{1, \ldots, K\}$$

The transformed mean and covariance are then approximated from these transformed points, keeping the weights unchanged, i.e.

$$\mathbf{m}_z = \sum_{i=1}^{K} w_i \mathcal{Z}_i$$
$$P_{zz} = \sum_{i=1}^{K} w_i (\mathcal{Z}_i - \mathbf{m}_z)(\mathcal{Z}_i - \mathbf{m}_z)^{\mathrm{T}}$$

Additionally, the cross-covariance between the input and the output can be computed, if desired, as

$$P_{xz} = \sum_{i=1}^{K} w_i (\mathcal{X}_i - \mathbf{m}_x)(\mathcal{Z}_i - \mathbf{m}_z)^{\mathrm{T}}$$

Any selection of sigma-points that exactly describes the input mean and covariance guarantees that the transformed mean and covariance is correctly calculated to second order [10].

One choice of sigma points is a set of $K = 2n + 1$ points that is chosen as

$$\begin{aligned}\mathcal{X}_0 &= \mathbf{m}_x \\ \mathcal{X}_i &= \mathbf{m}_x + \sqrt{n + \kappa}\, \mathbf{s}_{x,i} \\ \mathcal{X}_{i+n} &= \mathbf{m}_x - \sqrt{n + \kappa}\, \mathbf{s}_{x,i}\end{aligned}$$

with associated weights

$$\begin{aligned}w_0 &= \kappa/(n + \kappa) \\ w_i &= 1/2(n + \kappa) \\ w_{i+n} &= 1/2(n + \kappa)\end{aligned}$$

for $i \in \{1, \ldots, n\}$, where n is the dimension of the input $\mathbf{x}$, $\mathbf{s}_{x,i}$ is the i^{th} column of $\mathbf{S}_x$, $\mathbf{S}_x$ is a square-root factor of P_{xx} such that $P_x = \mathbf{S}_x \mathbf{S}_x^{\mathrm{T}}$, and κ is a tuning parameter of the UT. It is easily verified that this set of sigma-points matches the mean and covariance of $\mathbf{x}$; that is,

$$\mathbf{m}_x = \sum_{i=1}^{K} w_i \mathcal{X}_i$$
$$P_{xx} = \sum_{i=1}^{K} w_i (\mathcal{X}_i - \mathbf{m}_x)(\mathcal{X}_i - \mathbf{m}_x)^{\mathrm{T}}$$

B. Time Propagation

The UKF propagation step computes the *a priori* mean and covariance at time t_k (denoted $\hat{\mathbf{x}}_k^-$ and P_k^-, respectively) given the *a posteriori* mean and covariance at time t_{k-1} (denoted $\hat{\mathbf{x}}_{k-1}^+$ and P_{k-1}^+, respectively). An augmented state $\mathbf{z}_k$ is defined as

$$\mathbf{z}_{k-1}^{\mathrm{T}} = \begin{bmatrix} \mathbf{x}_{k-1}^{\mathrm{T}} & \mathbf{w}_{k-1}^{\mathrm{T}} \end{bmatrix}$$

Let the set of sigma points for the augmented state be denoted by the N values of $\mathcal{Z}_{i,k-1}$ and the associated weights by w_i where $i \in \{1, \ldots, N\}$ and $\sum_{i=1}^{N} w_i = 1$. These sigma points are generated from the augmented mean and covariance given by

$$\mathbf{m}_{k-1} = \begin{bmatrix} \hat{\mathbf{x}}_{k-1}^+ \\ \mathbf{0} \end{bmatrix} \qquad\qquad P_{k-1}^{\text{aug}} = \begin{bmatrix} P_{k-1}^+ & \mathbf{O} \\ \mathbf{O} & Q_{k-1} \end{bmatrix}$$

and each of the sigma points is partitioned as

$$\mathcal{Z}_{i,k-1}^{\mathrm{T}} = \begin{bmatrix} \mathcal{X}_{i,k-1}^{\mathrm{T}} & \mathcal{W}_{i,k-1}^{\mathrm{T}} \end{bmatrix}$$

The propagated sigma points are obtained via application of the nonlinear dynamical system, which gives

$$\mathcal{X}_{i,k} = \mathbf{f}(\mathcal{X}_{i,k-1}, \mathcal{W}_{i,k-1})$$

These transformed sigma-points are then used to approximate the nonlinear transformation of the mean and the covariance via

$$\hat{\mathbf{x}}_k^- = \sum_{i=1}^{N} w_i \mathcal{X}_{i,k} \tag{5}$$

$$P_k^- = \sum_{i=1}^{N} w_i (\mathcal{X}_{i,k} - \hat{\mathbf{x}}_k^-)(\mathcal{X}_{i,k} - \hat{\mathbf{x}}_k^-)^{\mathrm{T}} \tag{6}$$

While the effect of the process noise does not appear directly in these equations, it is captured through the propagated sigma points $\mathcal{X}_{i,k}$.

The attitude unscented estimator slightly modifies this algorithm by

1) Replaces the 4-dimensional portion of the sigma-points associated with the quaternion with a three-dimensional deviation between the i^{th} sigma-point quaternion and the 0^{th} sigma-point quaternion (hence $\mathcal{X}_0$ has a zero attitude deviation by construction)
2) Performs the algebraic weighted mean of the three-dimensional attitude deviations as if they were vectors
3) Calculates the propagated quaternion composing the 0^{th} sigma-point quaternion with the average three-dimensional deviation

While an advantage of the UT is that linearization is avoided by performing the algebraic average of three dimensional deviations, Ref. [7] effectively assumes they are vectors; hence, it assumes they are "small" and reside on the tangent space, which is an assumption consistent with the MEKF linearization.

C. Measurement Update

Using the propagated mean and covariance at time t_k, a new set of sigma points is created. Again, the first step is to define an augmented state

$$\tilde{\mathbf{z}}_k^{\mathrm{T}} = \begin{bmatrix} \mathbf{x}_k^{\mathrm{T}} & \mathbf{v}_k^{\mathrm{T}} \end{bmatrix}$$

and generate sigma points, along with their associated weights $\tilde{w}_i$, with the mean and covariance

$$\tilde{\mathbf{m}}_k = \begin{bmatrix} \hat{\mathbf{x}}_k^+ \\ \mathbf{0} \end{bmatrix} \qquad \tilde{P}_k^{\text{aug}} = \begin{bmatrix} P_k^- & O \\ O & R_k \end{bmatrix}$$

In this case, the set of sigma points for the augmented state is denoted by the $\tilde{N}$ values of $\tilde{\mathcal{Z}}_{i,k}$ where there is no restriction that $N = \tilde{N}$; that is, the update step may employ a different number of sigma points than the propagation step. Similarly, there is no requirement that the weights of the sigma points in the update step are the same as the weights of the sigma points in the propagation step. Each of the sigma points is partitioned as

$$\tilde{\mathcal{Z}}_{i,k}^{\mathrm{T}} = \begin{bmatrix} \mathcal{X}_{i,k}^{\mathrm{T}} & \mathcal{V}_{i,k}^{\mathrm{T}} \end{bmatrix}$$

and the measurement-transformed sigma points are given by

$$\mathcal{Y}_{i,k} = \mathbf{h}(\mathcal{X}_{i,k}, \mathcal{V}_{i,k})$$

The expected value of the measurement, the measurement covariance, and the cross-covariance are found in terms of the transformed sigma-points as

$$\hat{\mathbf{y}}_k^- = \sum_{i=1}^{\tilde{N}} \tilde{w}_i \mathcal{Y}_{i,k} \tag{7}$$

$$P_{yy,k} = \sum_{i=1}^{\tilde{N}} \tilde{w}_i (\mathcal{Y}_{i,k} - \hat{\mathbf{y}}_k^-)(\mathcal{Y}_{i,k} - \hat{\mathbf{y}}^-)^{\mathrm{T}} \tag{8}$$

$$P_{xy,k} = \sum_{i=1}^{\tilde{N}} \tilde{w}_i (\mathcal{X}_{i,k} - \hat{\mathbf{x}}_k^-)(\mathcal{Y}_{i,k} - \hat{\mathbf{y}}^-)^{\mathrm{T}} \tag{9}$$

The Kalman gain is $K_k = P_{xy,k}\ P_{yy,k}^{-1}$, and the associated updated state estimate and covariance are

$$\hat{\mathbf{x}}_k^+ = \hat{\mathbf{x}}_k^- + K_k\ (\mathbf{y}_k - \hat{\mathbf{y}}_k^-) \tag{10}$$

$$P_k^+ = P_k^- - K_k\ P_{yy,k}\ K_k^{\mathrm{T}} \tag{11}$$

Ref. [7] modifies the above algorithm by once again estimating three dimensional attitude deviations. Ref. [7] uses unit vector direction measurements and the classic additive measurement model:

$$\mathbf{y} = T\,\mathbf{r} + \mathbf{v} \tag{12}$$

The additive measurement model relies on linearization (for example, the large field-of-view model from Cheng *et al.* [13] linearizes around the actual measurement). Therefore, for coarse sensors, a multiplicative measurement model is more accurate in representing the actual error. Since both the measurement $\mathbf{y}$ and the reference vector $\mathbf{r}$ are of unit length, it follows from Eq. (12) that

$$\mathbf{y}^{\mathrm{T}}\mathbf{y} = 1 = \mathbf{r}^{\mathrm{T}}T^{\mathrm{T}}T\mathbf{r} + 2\mathbf{v}^{\mathrm{T}}T\mathbf{r} + \mathbf{v}^{\mathrm{T}}\mathbf{v} = 1 + 2\mathbf{v}^{\mathrm{T}}T\mathbf{r} + \mathbf{v}^{\mathrm{T}}\mathbf{v} \tag{13}$$

Taking expected values in Eq. (13) and using the fact that $\mathbf{r}$ is deterministic,

$$2\mathbf{r}^{\mathrm{T}}\ T\ \ \mathrm{E}\left\{\mathbf{v}\right\} = -\operatorname{trace}\mathrm{E}\left\{\mathbf{v}\mathbf{v}^{\mathrm{T}}\right\} \tag{14}$$

Eq. (14) implies that, for the classic additive measurement model of Eq. (12), the measurement noise is either zero mean with zero covariance or not zero mean, i.e. the measurement is biased.

The quaternion unscented estimator utilizes an additive residual $\mathbf{y}_k - \hat{\mathbf{y}}_k^-$, subtracting directions as if they were vectors. While this approach is shown to be optimal, contingent on the linearization assumption (i.e. it attains the Cramér-Rao lower bound [19]), no such guarantees exist when linearization is avoided with the UT.

D. Summary

Let's summarize a few remarks regarding the nature of the UKF algorithm. Firstly, the process of computing the propagated mean and covariance, Eqs. (5) and (6), relies on an averaging step and subtraction steps. Secondly, when considering the update stage of the UKF, vector subtraction is again utilized in Eq. (10). For situations in which unit vector measurements are to be processed, subtracting unit vectors will not yield a measurement residual that is also a unit vector. Therefore, when considering unit vector observations, the measurement update process of the UKF needs to be modified as well as the sample expectation calculations of Eqs. (7)–(9). All of these needed modifications may be grouped together as the removal of additivity within the UKF in favor of multiplicative steps.

IV. Alternative Attitude Filter

This section repeats the derivation of a fully multiplicative unscented quaternion estimator first presented in [11]. The filter uses twice the Gibbs vector to represent the attitude error (denoted as $\delta\mathbf{g}$). The attitude covariance is obtained from this three-dimensional quantity. In developing the filtering equations, the nature of rotations is preserved by never adding or subtracting three-dimensional attitude parameterizations nor unit-vector direction measurements. The linear minimum mean square error (LMMSE) estimate for nonlinear systems, of which the UKF is an approximation, seeks the estimate that minimizes the average of the square of the estimation error, which is usually defined as the Euclidean distance. The additive nature of the UKF described above is a direct result of the choice of the Euclidean distance. This algorithm minimizes the error (defined as twice the Gibbs vector), obtaining the minimum mean-square Gibbs error attitude estimate rather than minimum Euclidean error estimate.

The development that follows only includes the attitude in the state vector; adding other estimated quantities, such as a gyro bias, follows from a simple extension of the resulting equations.

A. Time Propagation

This portion of the algorithm calculates a propagated quaternion $\hat{\bar{\mathbf{q}}}$ and a propagated attitude error covariance matrix P^- that represent the estimates of how the attitude and its uncertainty evolve with time between measurement updates. During the propagation phase, a set of propagated sigma-point quaternions, $\hat{\bar{\mathbf{q}}}_j$, $j = 1, 2, \ldots, K$, is obtained following the same procedure of [7]; however, a different scheme is used to obtain the estimate. The desired estimate is the minimum mean-square error (MMSE) estimate. For a discrete random vector $\mathbf{X}$ with possible outcomes denoted by $\mathbf{x}_j$ and probability mass function p_j, the MMSE estimate $\hat{\mathbf{x}}$ minimizes

$$\hat{\mathbf{x}} = \min_{\mathbf{x}} \sum_j p_j \|\mathbf{x}_j - \mathbf{x}\|^2 \tag{15}$$

The solution of Eq. (15) is the mean of the random vector; that is,

$$\hat{\mathbf{x}} = \sum_j p_j \mathbf{x}_j \tag{16}$$

Prior attitude UKF implementations took the three-dimensional attitude parameterizations of the propagated sigma-point quaternions and performed their algebraic mean, effectively providing an estimate that minimizes the Euclidean distance between three-dimensional attitude errors. While this approach undoubtedly performs in a more than satisfactory fashion, it is more desirable to minimize an error defined as an attitude parameterization itself.

In this algorithm, the attitude estimation error is defined as twice the Gibbs vector ($\delta\mathbf{g}$, also known as Rodrigues parameters). The goal is to obtain the attitude MMSE estimate, which means minimizing the performance index

$$\hat{\bar{\mathbf{q}}} = \min_{\hat{\bar{\mathbf{q}}}} \sum_{j=1}^{K} w_j \|\delta\hat{\mathbf{g}}_j^-\|^2 \tag{17}$$

where

$$\bar{\mathbf{q}}(\delta\hat{\mathbf{g}}_j^-) = \delta\bar{\mathbf{q}} = \hat{\bar{\mathbf{q}}}_j \otimes \hat{\bar{\mathbf{q}}}^* \tag{18}$$

The asterisk represents the quaternion conjugate and the quaternion multiplication, represented by $\otimes$, composes quaternions in the same order as attitude matrices. Any other choice of attitude error representation will require a different performance index to be minimized and will produce a different estimate. Besides the aforementioned Euclidean distance, the easiest choice of attitude error would be to minimize the vector part of the quaternion error. Such a choice would produce an estimated quaternion with a known analytical solution obtained from solving a 4×4 eigenvalue problem [18]. The vector part of the quaternion is not a complete attitude parameterization; hence this algorithm minimizes an error that is physically a representation of attitude. Furthermore, the choice of the vector part of the quaternion as an error metric would produce undesirable effects during the measurement update phase of the algorithm, as detailed in the corresponding section of this paper.

The scaled Gibbs vector is given by

$$\delta\mathbf{g} = 2\delta\mathbf{q}_v / \delta q_s \tag{19}$$

where the subscripts v and s indicate the vector and scalar parts of the quaternion, respectively. The propagated quaternion estimate $\hat{\bar{\mathbf{q}}}$ is obtained by solving Eq. (17) numerically with a simple recursion. In all numerical simulations, a Newton-Raphson method is used and it always converges in very few iterations. The initial guess is chosen as the average quaternion in terms of minimizing the vector part of the quaternion error rather than the Gibbs vector, which is obtained by calculating the unit eigenvector corresponding to the maximum eigenvalue of

$$\mathbf{M} = 4 \sum_{j=1}^{K} \left(w_j \; \hat{\bar{\mathbf{q}}}_j \; \hat{\bar{\mathbf{q}}}_j^{\mathrm{T}} \right) - \mathbf{I}_{4\times4}$$

as shown in Ref. [18].

Particular care is also taken in computing the propagated attitude covariance. Once the estimated quaternion $\hat{\bar{\mathbf{q}}}$ is obtained from solving the eigenvalue problem, attitude deviations from the average quaternion are calculated for each sigma-point with Eq. (18), and the three-dimensional deviations $\delta\hat{\mathbf{g}}_j^-$ are then calculated with Eq. (19); no algebraic mean is ever performed. The propagated covariance is given by

$$P^- = \sum_{j=1}^{K} w_j \; \delta\hat{\mathbf{g}}_j^- \, (\delta\hat{\mathbf{g}}_j^-)^{\mathrm{T}} \tag{20}$$

Remark In order to reduce computations, it is possible to utilize the quaternion average proposed in [18] directly, as first done in [15], effectively producing the MMSE estimate where the error is defined as twice the vector part of the quaternion. In this work, however, we consistently define the error as the scaled Gibbs vector, which is a complete attitude parameterization.

B. Measurement Update

This proposed measurement update uses a multiplicative measurement model given by

$$\mathbf{y} = T(\boldsymbol{\eta}) \; T \; \mathbf{r} \tag{21}$$

where $\boldsymbol{\eta}$ is a three-dimensional representation of the attitude error, for example a rotation vector, $T(\boldsymbol{\eta})$ represents the direction cosine matrix parameterization of $\boldsymbol{\eta}$, T is the inertial-to-body coordinate transformation matrix, and $\mathbf{r}$ is the true direction in the inertial frame.

Using the multiplicative measurement model in the UKF overcomes the bias in the measurement error and allows an unbiased estimator to be obtained. Furthermore, one of the strengths of the UKF is that it avoids linearization around the mean; it is, therefore, more consistent to utilize a measurement model that also does not rely on linearization around a zero error as the QUEST measurement model or around the actual measurement as the large field-of-view model.

The second feature of the proposed update methodology is a multiplicative residual; unit vectors representing directions are not subtracted as if they were vectors in $\Re^3$. The attitude update is given by

$$\delta\hat{\mathbf{g}}^+ = \delta\hat{\mathbf{g}}^- + \mathbf{K}\boldsymbol{\epsilon} \tag{22}$$

where $\boldsymbol{\epsilon}$ is the multiplicative residual and, once again, the attitude error $\delta\mathbf{g}$ is twice the Gibbs vector defined as

$$\bar{\mathbf{q}}(\delta\mathbf{g}) = \bar{\mathbf{q}} \otimes \hat{\bar{\mathbf{q}}}^* \tag{23}$$

where $\bar{\mathbf{q}}$ is the true (unknown) inertial-to-body quaternion. From Eq. (23) it follows immediately that $\delta\hat{\mathbf{g}}^- = \mathbf{0}$; therefore, in fact, attitudes are never added together.

The residual expresses the "distance" between the actual measurement and the expected measurement; the greater this distance, the greater the update. To be consistent with our approach, we define the residual $\boldsymbol{\epsilon}$ as the scaled Gibbs vector that expresses the rotation to take $\hat{\mathbf{y}}$ into $\mathbf{y}$. There are infinite such rotations, so the minimum one is chosen, which is to say that we choose the Gibbs vector to be perpendicular to both $\hat{\mathbf{y}}$ and $\mathbf{y}$, which yields

$$\boldsymbol{\epsilon} = 2\frac{\hat{\mathbf{y}} \times \mathbf{y}}{1 + \hat{\mathbf{y}} \cdot \mathbf{y}} \tag{24}$$

where $\mathbf{y}$ is the unit vector measurement, which is one realization of the random vector $\mathbf{Y}$, and $\hat{\mathbf{y}}$ is the "average" measurement. Using the same logic employed before, the "average" measurement is the unit vector $\hat{\mathbf{y}}$ that minimizes the distance to all possible realizations of $\mathbf{Y}$. The distance is defined in terms of the Gibbs vector. Assuming a discrete distribution with possible outcomes denoted by $\mathbf{y}_j$ and probability mass function p_j

$$\hat{\mathbf{y}} = \min_{\hat{\mathbf{y}}} \sum_j p_j \|\boldsymbol{\epsilon}_j\|^2 = \min_{\hat{\mathbf{y}}} \sum_j p_j \; \frac{\|\hat{\mathbf{y}} \times \mathbf{y}_j\|^2}{(1 + \hat{\mathbf{y}} \cdot \mathbf{y}_j)^2} \qquad \text{subject to } \|\hat{\mathbf{y}}\| = 1$$

The minimizing value of $\hat{\mathbf{y}}$ is obtained numerically with a simple recursion.

Notice that Eq. (22) can be re-written as

$$\delta\hat{\mathbf{g}}^+ = \delta\hat{\mathbf{g}}^- + \mathbf{K}(\mathbf{z} - \hat{\mathbf{z}}) \tag{25}$$

where the auxiliary variable $\mathbf{z}$ is defined as $\mathbf{z} = 2(\hat{\mathbf{y}} \times \mathbf{y})/(1 + \hat{\mathbf{y}} \cdot \mathbf{y})$ and has zero mean, $\hat{\mathbf{z}} = \mathbf{0}$. Therefore this approach effectively seeks the MMSE estimate of $\mathbf{x}$ given the measurement $\mathbf{z}$. The proposed update is rewritten in the standard UKF form utilizing the auxiliary variable $\mathbf{z}$ and all the UKF properties still hold.

The sigma-points are obtained from the augmented covariance

$$P^{\text{aug}} = \begin{bmatrix} P^- & O \\ O & R \end{bmatrix} \tag{26}$$

where P^- is the *a priori* estimation error covariance and R is the measurement noise ($\boldsymbol{\eta}$) covariance. Because of the multiplicative measurement model of Eq. (21), R is chosen full-rank without any approximation. Linearized additive measurement models, on the other hand, possess a rank-deficient measurement error covariance. With the $n \times n$ matrix P^{aug} defined above, the $2n + 1$ sigma points are given by

$$\mathcal{X}_0 = \mathbf{0} \tag{27}$$

$$\mathcal{X}_i = \sqrt{(n+\kappa)\, P^{\text{aug}}_i} \tag{28}$$

$$\mathcal{X}_{i+n} = -\sqrt{(n+\kappa)\, P^{\text{aug}}_i} \tag{29}$$

where $i = 1, \ldots, n$ and $\sqrt{\mathbf{A}}_i$ is the i^{th} column of the matrix square root of $\mathbf{A}$. Along with the sigma-points, weights are chosen as

$$w_0 = \kappa/(n+\kappa) \qquad\qquad w_i = 0.5/(n+\kappa) \tag{30}$$

where κ is a design parameter of the UKF. Once the sigma points are obtained, they are transformed through the nonlinear measurement function as

$$\mathcal{Y}_i = \mathbf{h}(\mathcal{X}_i, \mathbf{r}, \hat{\bar{\mathbf{q}}}) \tag{31}$$

where

$$\mathbf{h}(\mathcal{X}_i, \mathbf{r}, \hat{\bar{\mathbf{q}}}) = T(\mathcal{N}_i)\, T(\delta\mathcal{G}_i)\, T(\hat{\bar{\mathbf{q}}})\, \mathbf{r} \tag{32}$$

In Eq. (32), $\delta\mathcal{G}_i$ and $\mathcal{N}_i$ are the elements that compose the input sigma points; that is,

$$\mathcal{X}_i^{\text{T}} = \begin{bmatrix} \delta\mathcal{G}_i^{\text{T}} & \mathcal{N}_i^{\text{T}} \end{bmatrix}$$

The mean and covariance of the transformed variables are found via

$$\hat{\mathbf{y}} = \min_{\boldsymbol{\xi}} \sum_{i=0}^{2n} w_i \frac{\boldsymbol{\xi} \times \mathcal{Y}_i}{1 + \boldsymbol{\xi} \cdot \mathcal{Y}_i}, \quad \|\boldsymbol{\xi}\| = 1 \qquad\qquad \mathcal{Z}_i = 2\frac{\hat{\mathbf{y}} \times \mathcal{Y}_i}{1 + \hat{\mathbf{y}} \cdot \mathcal{Y}_i} \tag{33}$$

$$P_{zz} = \sum_{i=0}^{2n} w_i\, \mathcal{Z}_i\, \mathcal{Z}_i^{\text{T}} \qquad\qquad P_{xz} = \sum_{i=0}^{2n} w_i\, \delta\mathcal{G}_i\, \mathcal{Z}_i^{\text{T}} \tag{34}$$

The updated state and covariance are obtained from Eq. (22) and

$$K = P_{xz} P_{zz}^{\dagger} \tag{35}$$

$$P^+ = P^- - K\, P_{zz}\, K^{\text{T}} \tag{36}$$

where the pseudoinverse $\dagger$ provides the optimal estimate given the singular covariance P_{zz}. Finally, the quaternion is updated as

$$\hat{\bar{\mathbf{q}}} \leftarrow \bar{\mathbf{q}}(\delta\hat{\mathbf{g}}^+) \otimes \hat{\bar{\mathbf{q}}} \tag{37}$$

Remark Once again, the system designer could choose to represent the error as twice the vector part of the quaternion rather than the scaled Gibbs vector, effectively minimizing $\sin^2 \theta/2$ rather than $\tan^2 \theta/2$, where θ is the Euler angle. For most, if not all, spacecraft applications, this alternative approach will work very well. However, the inherent constraint that each element of the vector part of the quaternion must be less than one can create problems in the presence of large attitude errors. A portion of the attitude sigma points are obtained as

$$\mathcal{X}_i = \sqrt{(n+\kappa)P_i^{\text{aug}}}$$

and there is no guarantee that for very uncertain systems $\sqrt{(n+\kappa)}$ will not scale the components of the attitude sigma points beyond unity. By choosing to represent attitude errors as Gibbs vectors, not only do we employ a full attitude parameterization, but we take advantage that the Gibbs attitude error is (almost) a one-to-one parameterization of attitude with the only exception/singularity being the 180 degree error. In the context of generating sigma points, that singularity is actually very helpful, because it prevents attitude sigma-points from wrapping around, allowing an extremely robust UKF design even for extremely large attitude uncertainties.

V. Numerical Comparison

The prior sections introduced the three algorithms to be compared and made clear some conceptual differences between the three:

1) The MEKF fully relies on linearization and the fact that small attitude is locally a vector space
2) The quaternion unscented estimator mixes features of the UT (which avoids linearization) with some assumptions that are valid only for vector spaces (attitude sigma-points are averaged, an additive unit vector measurement model is used, and an additive residual is used)
3) The multiplicative unscented Kalman filter removes all linearization and additivity assumptions for attitude states

This section numerically compares the performance of the three algorithms.

To demonstrate the validity of the proposed approach, we consider a satellite attitude tracking problem in which the orbit is perfectly known, but the attitude is not. The satellite is taken to be in near-geosynchronous orbit with Keplerian elements as shown in Table 1.

Table 1 Satellite Orbit

Type	Value	Units
Semi-Major axis	43000	km
Eccentricity	0.03	nd
Inclination	3	deg
RAAN	0	deg
Argument of Periapsis	0	deg
Mean Anomaly	0	deg

To generate a true attitude profile, we take the rotational dynamics to be

$$\dot{\bar{\mathbf{q}}} = \frac{1}{2}\bar{\omega} \otimes \bar{\mathbf{q}}$$
$$\dot{\omega} = J^{-1}\left(\sum \mathbf{m} - \omega \times J\omega\right)$$

where $\bar{\omega}$ is the pure quaternion formed from the angular velocity vector ω, J is the moment of inertia of the spacecraft, and $\sum \mathbf{m}$ represents the summation of all active moments in the body frame. The active moments are assumed to be zero in this work. The computation of the moment of inertia depends upon the mass distribution of the object. In this work, the object is assumed to be a hexagonal prism, as shown in Figure 1. This is an 8-plate model with the body-frame unit vectors defined by the unit vector triad $\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$. Additionally, the plate normal, denoted for the k^{th} plate by $\mathbf{u}^b_{n,k}$, is depicted in Figure 1. The area, A_k, and position from the object center, $\mathbf{r}^b_{p,k}$, of each plate are fully determined by specifying the side length, a, and the prism height, h. The distance of the side from the center, d, can be determined from the length of the side, a. The size parameters are chosen so as to represent a typical spacecraft size; the values

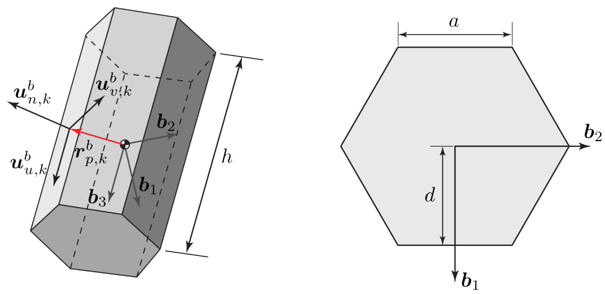

(a) 3D view (b) End view

Fig. 1 Satellite hexagonal prism flat plate model.

used, along with the total mass of the object, are presented in Table 2. Based upon the mass, side length, prism height, and the distance from the center to the side, the moment of inertia can be found to be a diagonal matrix of the form

$$J = \begin{bmatrix} I_{xx} & 0 & 0 \\ 0 & I_{yy} & 0 \\ 0 & 0 & I_{zz} \end{bmatrix}$$

where the elements of J are given by

$$I_{xx} = I_{yy} = m\left(\frac{a^2}{6} + \frac{d^2}{3} + \frac{h^2}{12}\right)$$
$$I_{zz} = m\left(\frac{a^2}{6} + \frac{d^2}{3}\right)$$

The resulting inertia values are also summarized in Table 2.

Table 2 Satellite Geometry and Characteristics

Type	Value	Units
Length of side	2	m
Height of side	4	m
Distance of side from center	1.7	m
Mass	2688	kg
I_{xx} and I_{yy} Inertia	8100	kg m^2
I_{zz} Inertia	4500	kg m^2

A. Nominal Test Case

We first present a nominal case for which we know the MEKF and the unscented quaternion estimator work properly. In this example, it is assumed that the initial attitude has a mean orientation given by the identity quaternion; that is, the mean quaternion represents a body frame that is exactly aligned to the inertial frame. Additionally, the initial mean angular velocity is taken to be zero. True values are generated by sampling a Gaussian error distribution with a standard deviation of 10° in attitude and 0.1 rev/day in angular velocity. The equations of motion specified previously are then applied to generate a true attitude and angular velocity profile.

The satellite is equipped with a three-axis rate-integrating gyro that provides incremental angular changes at 100 Hz. The gyro measurements are generated by integrating the true angular velocity signal at the 100 Hz frequency and then subjecting the true integrated signal to a zero-mean bias and a zero-mean white-noise sequence. The statistics of the gyro bias and noise are given in Table 3. In addition to the gyro, the satellite is equipped with a sun sensor and an Earth sensor operating at 1 Hz, which provide unit vector measurements that point to the sun and Earth, respectively. The pointing vectors are generated based on the specified (known) orbit and the uncertain attitude and then subjected to zero-mean white-noise sequences with standard deviations specified in Table 3.

Table 3 Sensor Specifications, Low-Noise Case

Type	1σ Error	Units
Gyro Noise	0.1	deg/$\sqrt{\text{s}}$
Initial Gyro Bias	1	deg/s
Sun Sensor Error	2	deg
Earth Sensor Error	5	deg

The attitude filter proposed in this work is then applied with a starting estimated quaternion equal to the identity quaternion and the estimated bias equal to zero. The initial attitude uncertainty and bias uncertainty that describe the elements of the initial covariance matrix are take to be 10° and 1°, 1σ, respectively. The UKF parameter is set to $\kappa = 3 - n$, and the filter's performance is shown in Figures 2 and 3. The gray line shows the estimation error while the black lines show the predicted 3σ error standard deviation.

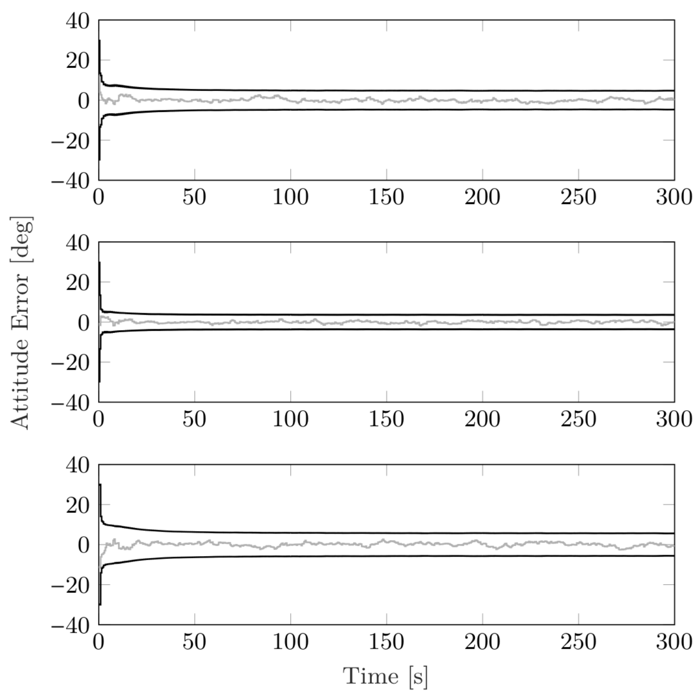

Fig. 2 Attitude estimation error for the low-noise test case.

Looking at Figures 2 and 3, it is clear that, in this single run of the proposed attitude filter, both the attitude and gyro bias are well-estimated. The attitude error converges to a steady-state value relatively quickly because the attitude is measured directly. The gyro bias, on the other hand, is estimated through its correlation with the attitude that is built during the propagation phase of the filter. It can be deduced from the figure that the gyro bias is estimated quite accurately, from the initial value of one degree (1σ), it reaches steady-state in about 150 seconds of simulation time and converges to less than 0.1 degrees (also a 1σ value). This example is to verify that the proposed algorithm works in a standard spacecraft scenario, a case for which we know other schemes, such as the MEKF and the UKF of Ref. [7], work well.

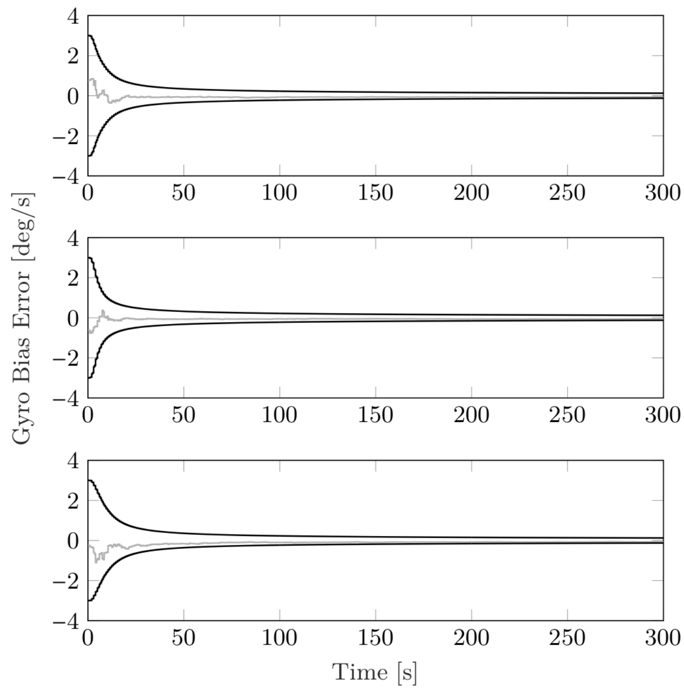

Fig. 3 Gyro bias estimation error for the low-noise test case.

B. High-Noise Test Case

In this example, we compare the algorithms when subjected to huge amounts of errors, much more than any existing and realistic sensor. It is assumed that the initial attitude has a mean orientation given by the identity quaternion and that the initial mean angular velocity is taken to be zero. True values are generated by sampling a Gaussian error distribution with a standard deviation of 50° in attitude and 0.1°/s in angular velocity. The sensors' error specifications are given in Table 4.

Table 4 Sensor Specifications

Type	1σ Error	Units
Gyro Noise	1	deg/$\sqrt{\text{s}}$
Gyro Bias	1	deg/s
Sun Sensor Error	50	deg
Earth Sensor Error	50	deg

All filters are initialized with a starting estimated quaternion equal to the identity quaternion and the estimated bias equal to zero and all three schemes use measurement and process noise variances equal to their true values in Table 4. The attitude estimation performance of the MEKF is established with 100 Monte Carlo runs and shown in Figure 4, where the blue line shows the 100 time histories of the estimation error while the red lines show the average of the 100 filter's predicted 3σ error standard deviation. It can be seen that the attitude error in the body-y and body-z directions is drifting and considerably outside of the filter's predictions. Figure 5 shows the performance of gyro bias estimation; this filter is not able to estimate the gyro bias.

The unscented quaternion estimator (UQE) is operated with $\kappa = 0$ to avoid any issues with a lack of positive definiteness in the filter's covariance matrix that sometimes plague UKF designs that choose $\kappa = 3 - n$ when $n > 3$ and large covariances. The resulting attitude estimation performance is summarized for 100 Monte Carlo runs in Figure 6, where the blue line shows the 100 time histories of the estimation error while the red lines show the average of the 100 filter's predicted 3σ error standard deviation. It can be seen that the filter is "smug," it under-represents the actual estimation error; the actual error remains outside the 3σ predicted uncertainty for long times. Figure 7 shows the

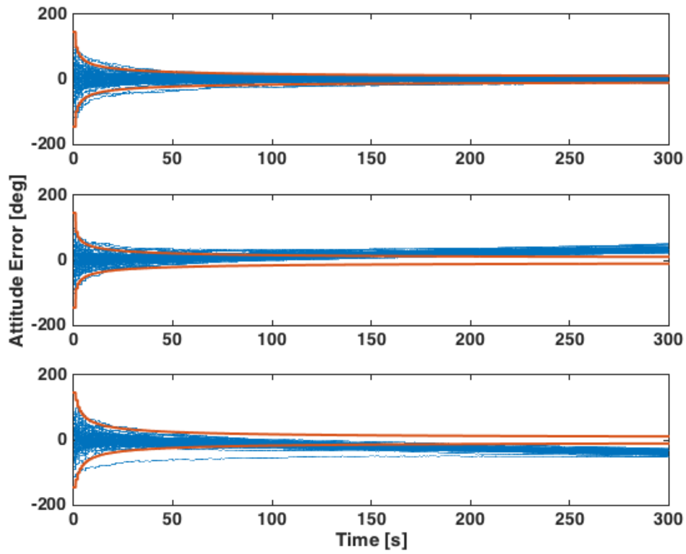

Fig. 4 Attitude estimation error (MEKF) - 100 MC runs

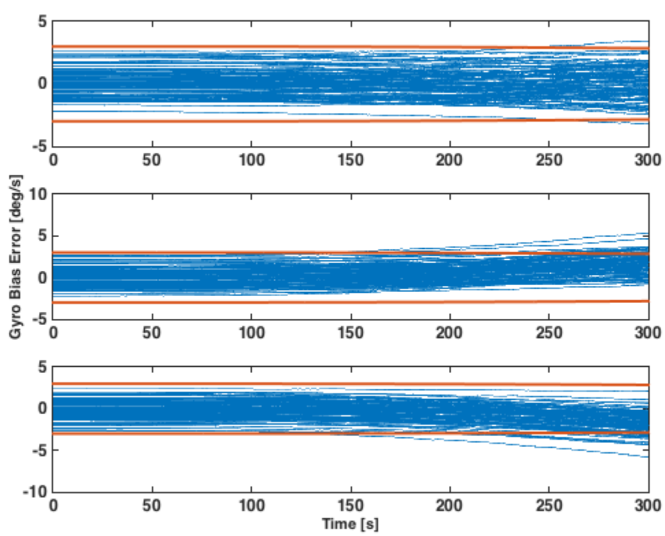

Fig. 5 Bias estimation error (MEKF) - 100 MC runs

performance of gyro bias estimation.

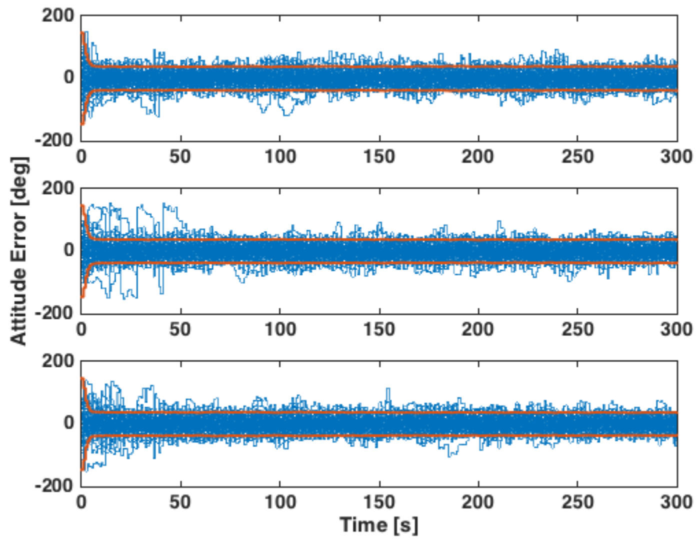

Fig. 6 Attitude estimation error (UQE) - 100 MC runs

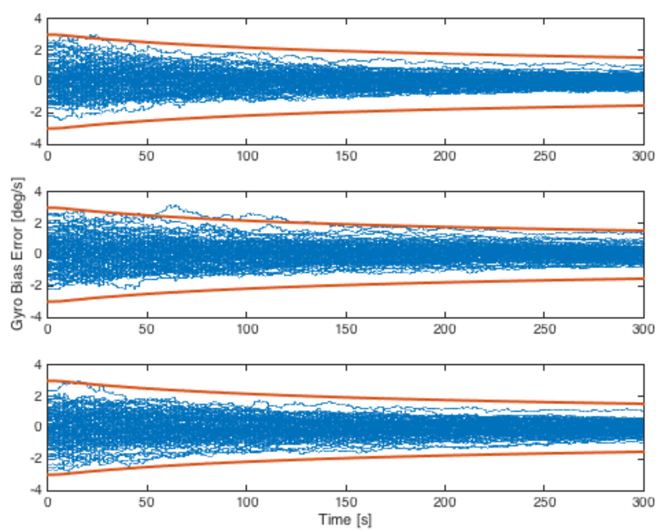

Fig. 7 Bias estimation error (UQE) - 100 MC runs

The fully multiplicative UKF (MUKF) has $\kappa = 0$ and the resulting attitude estimation performance is summarized for 100 Monte Carlo runs in Figure 8, where the blue line shows the 100 time histories of the estimation error while the red lines show the average of the 100 filter's predicted 3σ error standard deviation. It can be seen that the filter is capable of providing a consistent estimate. Figure 9 shows the performance of gyro bias estimation.

Discussion It is known that, in the presence of large initial uncertainties and accurate measurements, the EKF can diverge; this is due to the fact that the linearization of the nonlinear measurement function might not hold in the large domain of possible realizations of the state vector. In these situations, the EKF usually decreases its covariance too fast, and the actual estimation error lags behind and is not able to converge to its filter-predicted value. It is therefore

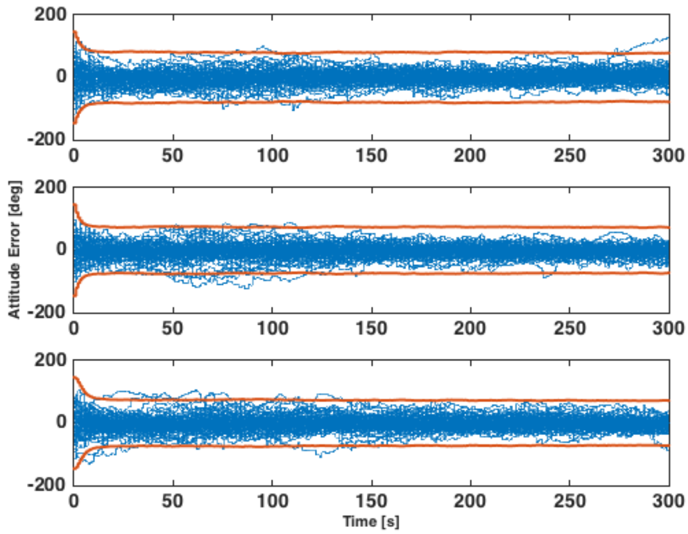

Fig. 8 Attitude estimation error (MUKF) - 100 MC runs

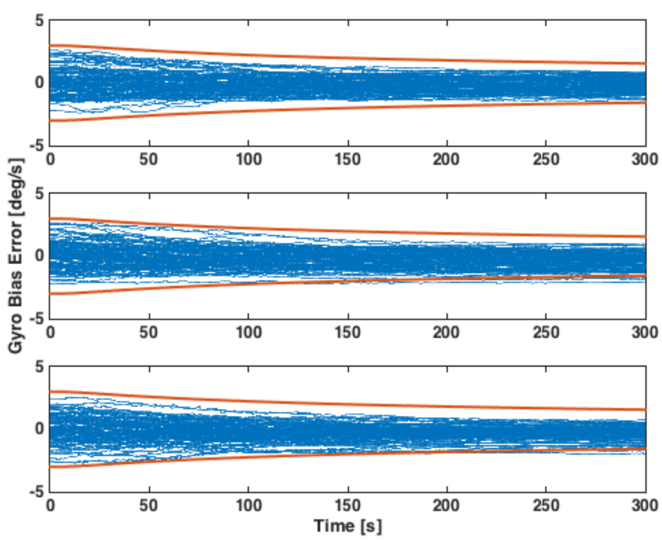

Fig. 9 Bias estimation error (MUKF) - 100 MC runs

known that the UKF can yield significantly better convergence and performance when starting with large uncertainty and utilizing precise measurements. Many EKF vs. UKF comparisons explore this type of scenario and conclude the superiority of the UKF.

Ref. [20] compares the attitude MEKF and UKF in another, quite unique, situation. They set the filter's value for the initial uncertainty of the gyro bias large (20 deg/hr, 1σ), but they make the actual initial error much smaller (5 deg/hr, 1σ). In this situation, the MEKF converges much faster than the UKF, and the authors conclude the MEKF is more robust to initial covariance discrepancies than the UKF is. If the actual initial bias error was also set to 20 deg/hr, 1σ, the MEKF covariance might shrink much faster than the actual error and provide a non-consistent estimate.

In this paper, we analyze a different scenario. Unlike [20], the covariances given to the filters match the true values and the filters are not tuned, i.e. the actual initial error covariance, process noise covariance, and measurement noise covariance are used by the algorithms without modifications. The initial error covariances are set to a large value, but unlike other comparisons that favor the UKF, the measurement error covariance is also kept large; otherwise, we know that the EKF could diverge. This approach gives a very fair comparison where none of the three algorithms are doomed to fail.

Under this high error situation we notice that:

1) The MEKF predicted covariance shrinks the fastest, and for the most part, the error initially follows it; however, the MEKF is not able to effectively estimate the gyro bias, which causes the attitude estimation error to drift away.
2) The unscented quaternion estimator, where additive residuals are used, produces an estimated covariance that is not consistent with the estimation error; the sample standard deviation obtained from 100 attitude error time series Monte Carlo runs is approximately twice the filter's prediction.
3) The recently proposed multiplicative unscented estimator produces an actual attitude estimator error similar to that of the unscented quaternion estimator (it seems a little less erratic towards the beginning of the simulation, but overall quite similar), but it has the advantage of producing a consistent covariance estimate. The sample error standard deviation from the Monte Carlo runs is slightly less than the filter's predicted standard deviation. The two UKF formulations estimate the gyro bias almost identically.

VI. Conclusions

This work analyzes the performance of two commonly used linear attitude filters, the multiplicative extended Kalman filter and the unscented quaternion estimator, with respect to a recently developed, fully multiplicative unscented Kalman filter for attitude estimation. The recently developed algorithms operates completely with multiplicative error, measurement, and residual models, enabling a consistent treatment of attitude and unit vectors throughout the filter. Simulations carried out on a nominal attitude estimation problem where line-of-sight measurements are processed demonstrate that the fully multiplicative unscented Kalman filter is capable of properly tracking the attitude and gyro bias when the uncertainties associate with the initial conditions and measurements are reasonable small. In a further test, it is shown that the new method also successfully tracks attitude and the gyro bias in high initial uncertainty and high measurement noise cases. The standard approach of the multiplicative extended Kalman filter, however, fails to estimate the gyro bias, leading to accumulation of attitude error. The unscented quaternion estimator produced good gyro bias estimates, but overly confident attitude estimates. Amongst the linear attitude estimators considered, the fully multiplicative unscented Kalman filter provided the most consistent estimates.

References

[1] Kuipers, J. B., *Quaternions and Rotation Sequences*, Princeton University Press, Princeton, NJ, 1999.

[2] Markley, F. L., "Attitude Estimation or Quaternion Estimation?" *The Journal of the Astronautical Sciences*, Vol. 52, No. 1/2, 2004, pp. 221–238.

[3] Lefferts, E. J., Markley, F. L., and Shuster, M. D., "Kalman Filtering for Spacecraft Attitude Estimation," *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 5, No. 5, 1982, pp. 417–429. doi:10.2514/6.1982-70.

[4] Bar-Itzhack, I. Y., and Oshman, Y., "Attitude Determination from Vector Observations: Quaternion Estimation," *IEEE Transaction on Aerospace and Electronic Systems*, Vol. 21, No. 1, 1985, pp. 128–135. doi: 10.1109/TAES.1985.310546 .

[5] Simon, D., and Chia, T. L., "Kalman Filtering with State Equality Constraints," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 38, No. 1, 2002, pp. 128–136. Doi: 10.1109/7.993234.

[6] Zanetti, R., Majji, M., Bishop, R. H., and Mortari, D., "Norm-Constrained Kalman Filtering," *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 5, 2009, pp. 1458–1465. doi: 10.2514/1.43119.

[7] Crassidis, J. L., and Markley, F. L., "Unscented Filtering for Spacecraft Attitude Estimation," *Journal of Guidance Control and Dynamics*, Vol. 26, No. 4, 2003, pp. 536–541. 10.2514/2.5102.

[8] Julier, S. J., Uhlmann, J. K., and Durrant-Whyte, H. F., "A new method for the nonlinear transformation of means and covariances in filters and estimators," *IEEE Transactions on Automatic Control*, Vol. 45, No. 3, 2000, pp. 477–482. doi: 10.1109/9.847726.

[9] Lefebvre, T., Bruyninckx, H., and Schutter, J. D., "Comment on "A New Method for the Nonlinear Transformation of Means and Covariances in Filters and Estimators"," *IEEE Transactions on Automatic Control*, Vol. 47, No. 8, 2002, pp. 1406–1408. Doi: 10.1109/TAC.2002.800742.

[10] Julier, S. J., and Uhlmann, J. K., "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, Vol. 92, No. 3, 2004, pp. 401–422. Doi: 10.1109/JPROC.2003.823141.

[11] Zanetti, R., and DeMars, K. J., "A New Multiplicative Unscented Kalman Filter for Attitude Estimation," *Journal of Guidance, Control, and Dynamics*, Accepted for publication.

[12] Shuster, M. D., "Kalman Filtering of Spacecraft Attitude and the QUEST Model," *The Journal of the Astronautical Sciences*, Vol. 38, No. 3, 1990, pp. 377–393.

[13] Cheng, Y., Crassidis, J. L., and Markley, F. L., "Attitude Estimation for Large Field-of-View Sensors," *The Journal of the Astronautical Sciences*, Vol. 54, No. 3 and 4, 2006, pp. 433–448. Doi:10.1007/BF03256499.

[14] Mortari, D., and Majji, M., "Multiplicative Measurement Model and Single-Point Attitude Estimation," *Proceedings of the F. Landis Markley Astronautics Symposium held June 30 – July 2, 2008, Cambridge, Maryland*, Advances in the Astronautical Sciences, Vol. 132, 2008, pp. 51–69. AAS 08-263.

[15] Zanetti, R., DeMars, K. J., and Mortari, D., "Novel Multiplicative Unscented Kalman Filter for Attitude Estimation," *Proceedings of the 2012 AAS/AIAA Space Flight Mechanics Meeting*, Advances in the Astronautical Sciences, Vol. 143, Charleston, South Carolina, 2012, pp. 337–348. AAS 12-125.

[16] Crassidis, J. L., Andrews, S. F., Markley, F. L., and Ha, K., "Contingency Designs for Attitude Determination of TRMM," *NASA/GSFC Flight Mechanics/Estimation Theory Symposium*, NASA/CP 3299, Greenbelt, MD, 1995, pp. 419–433.

[17] Zanetti, R., "A Multiplicative Residual Approach to Attitude Kalman Filtering with Unit-Vector Measurements," *The Journal of the Astronautical Sciences*, Vol. 57, No. 4, 2009, pp. 793–801. doi: 10.1007/BF03321530.

[18] Markley, F. L., Cheng, Y., Crassidis, J. L., and Oshman, Y., "Averaging Quaternions," *Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 4, 2007, pp. 1193–1197. Doi: 10.2514/1.28949.

[19] Hinks, J. C., and Crassidis, J. L., "Covariance Analysis of Maximum Likelihood Attitude Estimation," *The Journal of the Astronautical Sciences*, Vol. 60, No. 2, 2013, pp. 186–210. doi:10.1007/s40295-014-0028-7.

[20] Challa, M. S., Moore, J. G., and Rogers, D. J., "A Simple Attitude Unscented Kalman Filter: Theory and Evaluation in a Magnetometer-Only Spacecraft Scenario," *IEEE Access*, Vol. 4, 2016, pp. 1845 – 1858.