Fast Particle Filtering for Attitude and Angular-Rate Estimation from Vector Observations

Avishy Carmi* and Yaakov Oshman†
Technion—Israel Institute of Technology, 32000 Haifa, Israel

Extending and consolidating the recently introduced quaternion particle filter for a spacecraft's attitude estimation and its companion, the angular-rate particle filter, this paper presents a novel algorithm for the estimation of both a spacecraft's attitude and angular rate from vector observations. Belonging to the class of Monte Carlo sequential methods, the new estimator is a particle filter that uses approximate numerical representation techniques for performing the otherwise exact time propagation and measurement update of potentially non-Gaussian probability density functions in inherently nonlinear systems. The paper develops the filter and its implementation in the case of a low-Earth-orbit spacecraft, acquiring noisy Geomagnetic field measurements via a three-axis magnetometer. The new estimator copes with the curse of dimensionality related to the particle filtering technique by introducing innovative procedures that permit a significant reduction in the number of particles. This renders the new estimator computationally efficient and enables its implementation with a remarkably small number of particles (relative to the dimension of the state). The results of a simulation study demonstrate the viability and robustness of the new filter and its fast convergence rate.

I. Introduction

SPACECRAFT attitude is a critical piece of information in any space mission. In the last four decades, a great number of research works have been devoted to the problem of estimating the attitude of a spacecraft based on a sequence of noisy vector observations, resolved in the body-fixed coordinate system and in a reference system. The problem was first proposed in 1965 by Wahba [1] and inspired the development of the earliest so-called single-frame methods.

The source of body-fixed vector measurements can be a star-tracker, sun sensor, Earth sensor, and three-axis magnetometer (TAM). Whereas high-accuracy star-trackers are extremely expensive and sun sensors are useless during a solar eclipse [for low-Earth-orbit (LEO) satellites], the TAM is an integrated part of virtually any spacecraft, and its readings are available at any time.

The filtering-based methods that were developed in the 1980s embedded the attitude determination problem in the framework of stochastic filtering. The highlight of these methods is their ability to sequentially process vector measurements to yield the attitude at any time, by using an angular-rate information source as well. It quite naturally followed that the well-known Kalman filter (KF) was used as part of many attitude estimation algorithms.

The attitude estimation problem possesses an undesirable strong nonlinearity and may not be Gaussian either. It may also be subjected to a constraint related to the attitude mathematical representation. These distinctive features render conventional filtering algorithms (such as the KF) suboptimal at best.

The quaternion is probably the most popular rotation specifier in space applications. The main advantage of using the quaternion representation is that it is not singular for any rotation. Moreover, its kinematic equation is linear and the computation of the associated attitude matrix involves only algebraic expressions. However, the quaternion representation is not minimal because it is four dimensional. This leads to a normalization constraint that has to be addressed in filtering algorithms.

The quaternion was used in the framework of extended Kalman filtering (EKF) in [2] (using the so-called multiplicative approach) and [3] (using the additive approach). In a recent paper [4], an unscented Kalman filter (UKF) was proposed for the estimation of the rotation quaternion. The UKF does not require linearization. However, as a Kalman filter mechanization, it is sensitive to the statistical distribution of the stochastic processes driving the dynamic model: non-Gaussian distributions guarantee nonoptimality of the estimates.

A widely used angular-rate sensor onboard spacecraft is the rate-gyroscope triad, whose purpose is to provide three-axis rate information. Long experience has shown that rate gyros are failure prone. They tend to saturate during high-angular-rate scenarios, such as tumbling and initial attitude acquisition. Moreover, gyros may not be suitable for low-cost satellites due to price, power consumption, weight, and volume considerations. This leads to the requirement of reliable, gyroless attitude estimation schemes that will provide backup capability for spacecraft that use rate gyros and affordable solutions for low-cost gyroless satellites.

Several methods were proposed for combined attitude and angular-rate filtering in the absence of rate sensors. Most of these methods adopt the EKF mechanization; therefore, they also suffer from the same deficiencies imposed by applying linearization and assuming Gaussianity. Moreover, due to the higher nonlinearity involved in the augmented models, which now account for the angular-rate dynamics as well, linearization renders these methods much more sensitive to initial conditions.

In [5], high-bandwidth star-tracker vector observations were used by an EKF-based algorithm to estimate both the attitude quaternion and the angular velocity of the spacecraft in a gyroless attitude determination and control setting. Challa et al. [6,7] and Natanson et al. [8] proposed an attitude and attitude-rate estimator, which uses temporal derivatives of vector measurements and dynamically propagates the angular velocity estimates using the nonlinear Euler equation. The technique of Bar-Itzhack et al. [9], on the other hand, consists of an unusual KF formulation that renders the models used

*Doctoral Student, Department of Aerospace Engineering; avishy@aerodyne.technion.ac.il.
†Professor, Department of Aerospace Engineering and Member, Technion's Asher Space Research Institute; also, Holder of the Louis and Helen Rogow Chair in Aeronautical Engineering; yaakov.oshman@technion.ac.il. Fellow AIAA.

by the filter state dependent. This special pseudolinear Kalman filter (called PSELIKA by its authors) possesses the advantage that it is less likely to diverge; however, its estimation accuracy is likely to be lower than that of the conventional EKF.

Recently, a new method using the particle filtering (PF) technique has been proposed for a spacecraft's attitude estimation from vector observations [10]. The PF technique was further employed in [11] for the exclusive estimation of angular rate independent of any attitude information.

Also known as sequential Monte Carlo methods, particle filters implement recursive Bayesian models using simulation-based methods [12]. Avoiding the underlying assumptions of the Kalman filter, namely, that the state space is linear and Gaussian, these rather general and flexible methods enable solving for the posterior probability distributions of the unknown variables (on which all inference on these variables is based) within a Bayesian framework, exploiting the dramatic recent increase in computing power. It should be emphasized that PFs are not just smart implementations of the Kalman filter or its nonlinear variants/extensions; rather, they are entirely different algorithms that lead to entirely different solutions to the nonlinear, non-Gaussian filtering problem. Contrary to the Kalman filter extensions, the solutions obtained using PF algorithms are approximations to the optimal (in the Bayesian sense) solutions, which can be made arbitrarily close to the exact solutions by increasing the number of particles involved in the computation and thereby increasing the computation workload.

This paper presents a unification of the quaternion particle filter (QPF) of [10] with the angular-rate PF of [11]. The new PF is implemented for an augmented seven-element state vector that accounts for four quaternion and three angular-rate components. A straightforward implementation of such PF would require a large number of particles to work properly. Thus, the main contribution of this work is in deriving a computationally efficient high-dimensional PF algorithm for estimating both the attitude quaternion and the angular rate. The new algorithm implements several innovative techniques for significantly reducing the number of particles used in the approximations.

Estimating the attitude quaternion and angular rate using a PF provides three major advantages relative to existing estimation methods. First, the estimator works directly with the quaternion (i.e., its particles are all attitude quaternions). This inherently renders its resulting estimate an attitude quaternion, naturally satisfying the norm constraint (requiring no special procedures for this purpose). Second, because the PF algorithm assumes nothing about the noise distributions, the resulting algorithm can work with any noise distribution associated with the particular sensors involved. In contradistinction, the UKF assumes Gaussian distributions of the driving noise processes, which does not always hold true. Finally, the PF is easy to implement and is insensitive to the initial conditions and nonlinearities involved.

In general, gyroless setting inevitably requires the use of Euler's equation of rigid body motion in filtering algorithms. As was pointed out by both Carmi et al. [11] and Tortora et al. [13], this renders the time-propagation stage highly sensitive to a spacecraft's inertia modeling imperfections. Consequently, recent tendency is to estimate the inertia tensor as well. As far as particle filtering is concerned, simply augmenting the state by the inertia components results in a substantial increase in the algorithm's computational complexity. Nevertheless, on a principle level, this work demonstrates the feasibility of inertia estimation using an application of the conventional state-augmentation technique.

The remainder of this paper is organized as follows. The next section outlines the mathematical model of the quaternion/angular-rate estimation problem. Section III provides a detailed development of the computationally efficient attitude and angular-rate particle filter. As part of the development, techniques for significantly reducing the computational resources used by the PF are presented. Section IV presents the results of a simulation study that was carried out to assess the performance of the new algorithm. Concluding remarks are offered in the last section.

II. Mathematical Model

In this section, the problem of combined attitude and angular-rate estimation from vector observations is mathematically defined.

A. Observation Model

Let r_k and y_k be a pair of corresponding vector measurements acquired at time k in the two Cartesian coordinate systems $\mathcal{R}$ and $\mathcal{B}$, respectively. Let A_k be the rotation matrix (also known as the attitude matrix or the direction cosine matrix) that brings the axes of $\mathcal{R}$ onto the axes of $\mathcal{B}$ at time k. In general, the reference vector r_k is known exactly, whereas the body vector $\boldsymbol{b}_{y\,k}$ is measured. This results in the following attitude measurement model:

$$y_k = A_k r_k + n_k \qquad (1)$$

where n_k is the measurement noise process, with a known probability density function (PDF), denoted as $n_k \sim p_{n_k}$.

B. Process Model

1. Angular-Rate Process Model

Let $\{\omega_k\}_{k=0}^{\infty}$, $\omega \in \mathbb{R}^3$ be the angular velocity process of the body-frame Cartesian coordinate system $\mathcal{B}$ with respect to some inertial reference frame coordinate system $\mathcal{R}$, resolved in $\mathcal{B}$, at times $k = 0, 1, 2, \ldots$. Representing the angular rate of a spacecraft, this process is the discrete-time equivalent of a continuous-time stochastic process described via Euler's equation. Using the common formal engineering notation, this stochastic equation is written as

$$\dot{\omega}(t) = J^{-1}[-\omega(t) \times J\omega(t)] + v(t) \qquad (2)$$

where J denotes the spacecraft's tensor of inertia. Given some initial distribution $\omega_0 \sim p_{\omega_0}$, Eq. (2) is driven by the process noise $v(t)$, representing the external disturbance torques.

2. Quaternion Process Model

The discrete-time quaternion stochastic process satisfies the recurrence equation

$$\mathbf{q}_{k+1} = \Phi_k \mathbf{q}_k \qquad (3)$$

where $\{\mathbf{q}_k\}_{k=0}^{\infty}$ denotes the quaternion of rotation from a given reference frame $\mathcal{R}$ onto the body frame $\mathcal{B}$ at times $k = 0, 1, 2, \ldots$, with some initial PDF $\mathbf{q}_0 \sim p_{\mathbf{q}_0}$. The quaternion process takes its values on the unit three-sphere $\mathbb{S}^3$ and is constructed from vector and scalar parts, respectively:

$$\mathbf{q}_k = \begin{bmatrix} \varrho_k^T & \mathrm{q}_{4_k} \end{bmatrix}^T \qquad (4)$$

The orthogonal transition matrix Φ_k is expressed using ω_k. Assuming that ω_k is constant during the short sampling time interval $\Delta t \triangleq t_{k+1} - t_k$ yields

$$\Phi_k \triangleq \Phi(\omega_k) = \exp\left(\frac{1}{2}\Psi(\omega_k)\Delta t\right), \qquad \Psi(\omega_k) \triangleq \begin{bmatrix} -[\omega_k\times] & \omega_k \\ -\omega_k^T & 0 \end{bmatrix} \tag{5}$$

where $[\omega_k\times]$ denotes the cross-product matrix associated with the vector ω_k.

The relation between the process and observations is established by expressing the attitude matrix as a quadratic function of $\mathbf{q}$, that is,

$$A_k = A(\mathbf{q}_k) = \left[(\mathrm{q}_{4_k})^2 - \varrho_k^T \varrho_k\right] I_{3\times3} + 2\varrho_k \varrho_k^T - 2\mathrm{q}_{4_k}[\varrho_k\times] \qquad (6)$$

III. Fast High-Dimensional Attitude and Angular-Rate Particle Filtering

The fast quaternion and angular-rate PF is derived in this section. Initially, the plain PF algorithm is presented, followed by the incorporation of several techniques for reducing its computational complexity. In this PF algorithm, each particle is a state vector

composed of a unit norm quaternion and a three-element angular-rate vector. The filter is aimed at approximating the joint filtering PDF of both the attitude and the angular rate conditioned upon the observation time history. At this stage, it is assumed that the spacecraft's inertia tensor is perfectly known. The case of imperfectly modeled inertia is discussed thereafter.

A. Particle Filters

Particle filters are a set of simulation-based methods using a Bayesian framework for sequentially approximating the posterior PDF of arbitrary nonlinear non-Gaussian systems. Fundamentally, based upon both the strong law of large numbers and a statistical sampling method known as importance sampling, particle filters exploit the recent increase in computational resources to yield superior performance compared with the conventional nonlinear filtering methods. Using these methods, a discrete representation of the posterior PDF is maintained up to a normalizing constant via a finite set of samples (particles). The PDF approximation is propagated using the Chapman–Kolmogorov equation and updated using Bayes rule. Practical implementation of a PF algorithm consistently involves adaptation of particles and their associated importance weights. The reader is referred to [10,12] for either a brief or detailed development, respectively, of a generic PF algorithm.

B. Plain Particle Filtering Algorithm

The plain attitude and angular-rate PF borrows its structure and components from both the QPF and the angular-rate PF (OPF) algorithms, presented in [10,11], respectively. To maintain consistency with previous work, the notational conventions of both these references are adopted. Accordingly, $\mathcal{Y}^k = \{y_1, \ldots, y_k\}$ and $Y^k = \{Y_1, \ldots, Y_k\}$ denote a set of vector observations up to time k and its realization, respectively. The quaternion and the angular-rate realizations at time k are denoted as q_k and Ω_k, respectively. Likewise, the ith particle and its associated normalized importance weight at time k are denoted by $[q_k(i)^T, \Omega_k(i)^T]^T$ and $\tilde{w}_k(i)$, respectively. In what follows, the term quaternion/angular-rate part refers to the corresponding part of the joint state vector.

In the plain PF algorithm, the measurement update stage as well as the quaternion part filtering and particles' evolution stages are similar to those used in [10] (Secs. IV.B, IV.C. and IV.E, respectively). The angular-rate part is propagated and filtered using the approach of [11] (Secs. IV.C and IV.D, respectively). The algorithmic details of these stages are not repeated herein for conciseness. The main difference underlying the plain PF algorithm, when compared with either of the algorithms of [10,11] lies in its particles' maintenance stage. As part of this procedure, a new high-dimensional regularization scheme is incorporated into the algorithm, taking into account the statistical dependency between the quaternion and the angular-rate parts. This two-step procedure is detailed next.

1. Regularization Intensity

The regularization intensity (RI) is a measure of the contaminating noise used for diversifying the particle set. Considering the unique case of a partially constrained state vector (the four components of the quaternion), the computation of the RI differs from the one introduced in [10]. In this work, the RI measure is computed as the sample covariance of an equivalent six-dimensional unconstrained state, taking into account the cross-correlation terms between the attitude and the angular rate.

During this stage, the four-dimensional quaternion part of the state is replaced by the corresponding three-dimensional generalized Rodrigues parameters (GRP) vector. Given a quaternion realization q_k, the associated GRP vector, denoted by ξ_k, is computed as

$$\xi_k = f \frac{\rho_k}{a + q_{4_k}} \tag{7}$$

where ρ_k is the vector part of the quaternion realization (i.e., the realization of ϱ_k), $a \in [0, 1]$, and f is a scale factor. In this work, the values of $a = 1$, $f = 2(1 + a) = 4$ are chosen, so that ξ_k is equal to

the Euler parameters for small angles. An empirical measure of the covariance of $[q_k^T, \omega_k^T]^T$ conditioned upon the observation history, $\mathcal{Y}^k$, is then obtained as

$$\hat{P} = \sum_{i=1}^{N} \tilde{w}_k(i) \left\{ \left[\xi_k(i)^T - \hat{\xi}_k^T, \Omega_k(i)^T - \hat{\Omega}_k^T \right]^T \right.$$
$$\left. \times \left[\xi_k(i)^T - \hat{\xi}_k^T, \Omega_k(i)^T - \hat{\Omega}_k^T \right] \right\} \tag{8}$$

where $\xi_k(i)$ denotes the GRP associated with the quaternion part of the ith particle $q_k(i)$, and $\hat{\xi}_k$ denotes the GRP associated with the filtered quaternion part at time k.

Remark 1. When using Eq. (8), one should note that, although both quaternions $+q$ and $-q$ represent a single unique attitude, the corresponding modified Rodrigues parameters have two different values. Thus, one must ensure that all particles are composed of quaternion parts having the same scalar-part sign.

2. Regularization Scheme

Let $\{[q_k(i)^T, \Omega_k(i)^T]^T\}_{i=1}^N$ be the stock of N state samples at time k.

The number of offspring for each particle, $N_k(i)$, is determined following the approach of [10]. The particle offspring are then computed in the following manner.

First, a set of vectors is sampled from a six-dimensional zero-mean unit covariance Gaussian (or some other) kernel, denoted by $\mathcal{K}$, to obtain

$$[\delta\xi_o(j)^T, \delta\Omega_o(j)^T]^T \sim \mathcal{K}(0, I_{6\times6}), \qquad j = 1, 2, \ldots, N_k(i)/2 \tag{9}$$

where $\delta\xi_o(j) \in \mathbb{R}^{3\times1}$, $\delta\Omega_o(j) \in \mathbb{R}^{3\times1}$. When $N_k(i)$ is odd, one of the values of the vector $[\delta\xi_o(j)^T, \delta\Omega_o(j)^T]^T$ is taken to be $0_{6\times1}$ (ensuring a symmetric set around $[q_k(i)^T, \Omega_k(i)^T]^T$). The next step consists of rescaling and rotating the six-dimensional space according to the previously obtained RI measure. Thus, each vector is rescaled and rotated according to

$$[\delta\xi_k(j)^T, \delta\Omega_k(j)^T]^T = h\hat{P}^{1/2}[\delta\xi_o(j)^T, \delta\Omega_o(j)^T]^T \tag{10a}$$

$$[\delta\xi_k(j + N_k(i)/2)^T, \delta\Omega_k(j + N_k(i)/2)^T]^T$$
$$= -h\hat{P}^{1/2}[\delta\xi_o(j)^T, \delta\Omega_o(j)^T]^T \tag{10b}$$

for $j = 1, 2, \ldots, N_k(i)/2$, where h is some predetermined bandwidth, and $\hat{P}^{1/2}$ denotes the matrix square root of $\hat{P}$. In this work, the bandwidth h is set as suggested in [12] (p. 253), that is,

$$h = [4/(N(n + 2))]^{\frac{1}{n+4}} \tag{11}$$

with $n = 7$, corresponding to the state dimension. Notice that Eq. (10) produces a symmetric set of $N_k(i)$ vectors, $\{[\delta\xi_k(l)^T, \delta\Omega_k(l)^T]^T\}_{l=1}^{N_k(i)}$. The offspring of the ith particle $[q_k(i)^T, \Omega_k(i)^T]^T$, denoted as $\{[\bar{q}_k(l)^T, \bar{\Omega}_k(l)^T]^T\}_{l=1}^{N_k(i)}$, are then computed using the inverse transformation of Eq. (7) (with $a = 1$, $f = 4$), as

$$\begin{bmatrix} \bar{q}_k(l) \\ \bar{\Omega}_k(l) \end{bmatrix} = \begin{bmatrix} \frac{1}{4}(1 + \bar{q}_{4_k}(l))[\xi_k(i) + \delta\xi_k(l)] \\ \bar{q}_{4_k}(l) \\ \Omega_k(i) + \delta\Omega_k(l) \end{bmatrix} \tag{12a}$$
$$l = 1, \ldots, N_k(i)$$

where

$$\bar{q}_{4_k}(l) = \frac{16 - \|\xi_k(i) + \delta\xi_k(l)\|^2}{16 + \|\xi_k(i) + \delta\xi_k(l)\|^2}, \qquad l = 1, \ldots, N_k(i) \tag{12b}$$

The new set of particles thus obtained are assumed to be equally weighted (i.e., all importance weights are set to $1/N$) or, alternatively, can be reweighted using the likelihood PDF associated with the last vector measurement (see [10]).

A single iteration of the plain quaternion and angular-rate PF is illustrated in Fig. 1.
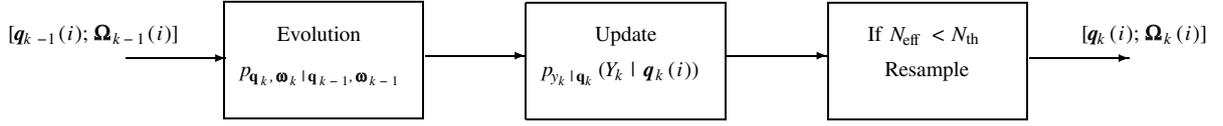
Fig. 1 Quaternion and angular-rate PF.

C. Computationally Efficient High-Dimensional Particle Filtering

The plain PF processes a seven-element state vector. Practical implementation of such a filter requires a large number of particles, a phenomenon known as the curse of dimensionality. In fact, experience shows that, when applied to the combined attitude and rate problem, the plain PF requires at least 30,000 particles to attain an acceptable level of performance.

In this subsection, three techniques are proposed for significantly reducing the number of particles used by the PF algorithm. Consisting of an efficient initialization procedure that exploits the unique characteristics of the problem, the first technique aims at reducing the PF search space (sample space). The other techniques deal with another aspect of the dimensionality problem.

1. Particle Filtering Initialization

The size of the PF search space can be easily reduced by concentrating particles in high-likelihood areas only. Identifying such areas may consist of using smart initialization procedures, such as those implemented in [10,11]. In these methods, the initial set of particles was generated based upon the first measurement. The rationale behind this type of initialization is simple: instead of starting with no knowledge at all, namely, with a uniformly distributed set of particles, say, p_{x_0}, the first measurement can be used to generate a set of particles for approximating $p_{\mathbf{q}_0, \omega_0 | y_0}$. Assuming that the vectors $\mathbf{q}_0$ and ω_0 are independent and uniformly distributed, it follows from Bayes theorem that

$$p_{\mathbf{q}_0, \omega_0 | y_0}(\boldsymbol{q}_0(i), \boldsymbol{\Omega}_0(i) | \boldsymbol{Y}_0) \propto p_{y_0 | \mathbf{q}_0}(\boldsymbol{Y}_0 | \boldsymbol{q}_0(i)) p_{\mathbf{q}_0}(\boldsymbol{q}_0(i)) p_{\omega_0}(\boldsymbol{\Omega}_0(i))$$

$$\propto p_{y_0 | \mathbf{q}_0}(\boldsymbol{Y}_0 | \boldsymbol{q}_0(i)) \qquad (13)$$

Equation (13) implies that the first vector measurement $\boldsymbol{Y}_0$ can be used to generate the initial set of quaternions independently of the angular rate. In much the same way, the initial set of the angular-rate particles can be generated independently of the attitude. At least two vector measurements are needed to generate such a set. Detailed algorithms for producing the initial set of quaternion and angular-rate particles can be found in [10,11], respectively.

The two aforementioned algorithms are applied to the first two vector measurements. The angular-rate particle-set-generation procedure is based on the body-frame vectors of the two consecutive measurements, whereas the quaternion particles are extracted using the second vector measurement only. The resulting quaternion and angular-rate particles populate the more likely regions of $p_{\mathbf{y}_1 | \mathbf{q}_1}$ and $p_{\mathcal{y}^1 | \omega_1}$, respectively. In the last stage of the initialization, the two sets of particles, namely, the quaternion and the angular-rate particles, are combined to yield a joint initial set.

Let $\{\boldsymbol{q}_1(j)\}_{j=1}^{N_{\mathbf{q}}}$ and $\{\boldsymbol{\Omega}_1(l)\}_{l=1}^{N_\omega}$ be the initial sets of quaternion and angular-rate vector particles, respectively. Recognizing that the elements of each of these sets are nearly equally weighted (due to the fact that a single vector observation determines the rotation up to 1 degree of freedom [10,11]), the initial set of particles for approximating the joint PDF $p_{\mathbf{q}_1, \omega_1 | \mathcal{y}^1}$ is taken as the Cartesian product $\{\boldsymbol{q}_1(j)\}_{j=1}^{N_{\mathbf{q}}} \times \{\boldsymbol{\Omega}_1(l)\}_{l=1}^{N_\omega}$, resulting in $N = N_{\mathbf{q}} N_\omega$ particles, $\{[\boldsymbol{q}_1(j)^T, \boldsymbol{\Omega}_1(l)^T]^T, \ j = 1, \ldots, N_{\mathbf{q}}, \ l = 1, \ldots, N_\omega\}$.

Further improvement of the initial particle set is achieved by applying the angular-rate PF of [11] to a batch of measurements $\mathcal{y}^{k_1}$, and then selecting N_ω initial angular-rate particles out of its population. The initial set of quaternion particles, $\{\boldsymbol{q}_{k_1}(j)\}_{j=1}^{N_{\mathbf{q}}}$, is produced as before, this time using the vector observation $\boldsymbol{y}_{k_1}$. In this case, the joint initial particle set is taken as $\{[\boldsymbol{q}_{k_1}(j)^T, \boldsymbol{\Omega}_{k_1}(l)^T]^T, \ j = 1, \ldots, N_{\mathbf{q}}, \ l = 1, \ldots, N_\omega\}$. This procedure is equivalent to approximating the joint PDF at time k_1 as

$$p_{\mathbf{q}_{k_1}, \omega_{k_1} | \mathcal{y}^{k_1}}(\boldsymbol{q}_{k_1}(i), \boldsymbol{\Omega}_{k_1}(i) | \boldsymbol{Y}^{k_1})$$

$$\approx c \, p_{\mathbf{y}_{k_1} | \mathbf{q}_{k_1}}(\boldsymbol{Y}_{k_1} | \boldsymbol{q}_{k_1}(i)) p_{\omega_{k_1} | \mathcal{y}^{k_1}}(\boldsymbol{\Omega}_{k_1}(i) | \boldsymbol{Y}^{k_1}) \qquad (14)$$

where $c = p_{\mathbf{q}_{k_1}}(\boldsymbol{q}_{k_1}(i)) / p_{\mathbf{y}_{k_1}}(\boldsymbol{Y}_{k_1})$ is a normalization constant (assuming $\mathbf{q}_{k_1}$ is uniformly distributed). When using this batch initialization technique, the plain PF algorithm is executed subsequent to the acquisition of k_1 measurements.

Remark 2. As is well known to estimation practitioners, in the EKF and UKF algorithms, good initialization is vital for the ensuing filter's convergence and accuracy, because of the approximations involved in both algorithms. Thus, both algorithms cannot be used in "lost in space" applications. In contradistinction, in the PF presented herein, the main purpose of the batch initialization procedure, employed on the angular-rate channel, is not to improve the quality of the initialization (and, thus, the ensuing filter's performance). Rather, the purpose of the initialization procedure is to reduce the number of particles used at the initial stage of the filter to achieve increased computational efficiency. Whereas in both the EKF and UKF poor initialization may result in unavoidable divergence, in the PF the initialization procedure can be skipped altogether without sacrificing the filter's estimation performance simply by increasing the number of particles used during the initial stages. Thus, the smart initialization procedure effectively reduces the computational load (related to the number of particles used), but is, by no means, vital to the filter's ensuing performance.

2. Search Space Local Traps

The PF algorithm maintains a discrete approximation of the posterior PDF via a finite number of samples. Because of the finiteness of the particle set, the actual search space (sample space), considered by the PF, contains "holes." These holes are regions of the search space that are, incorrectly, regarded by the PF to have zero probabilities. As long as relatively small probability regions of the search space are considered holes, the PF algorithm works properly. Unfortunately, the converse situation may occur also. In such cases, the particles are ultimately trapped in small probability regions of the space and the PF fails to approximate the posterior PDF.

Intuitively, it can be thought that the size of a hole is related to the number of particles used for the approximation, their locations in the search space, and the dimension of the entire space. It easily follows that state augmentation, which implies a larger search space, results in a higher probability for the PF algorithm to be trapped in a hole.

To understand this phenomenon, the particle maintenance phase is examined. As part of this stage, the particle trajectories are manipulated via resampling and regularization in accordance with their importance weights.

1) *Resampling*: This procedure causes a deterioration of the approximation by damaging the particle set diversity, consequently creating holes. During resampling, lightly weighted particles are rapidly discarded, thus driving the entire particle population to the vicinity of highly weighted locations in the search space but leaving other regions unpopulated.

2) *Regularization*: As opposed to resampling, this procedure is aimed at diversifying the particle set by deliberate contamination; hence, it can be regarded as a "hole filler." The intensity of contamination, which controls the expansion of the new particle set, is related to the sample covariance of the particle set or to some other similar measure.

3) *Importance Sampling*: Setting the importance density as close as possible to the posterior PDF tends to diminish the phenomenon.

In that case, the PF can better evaluate the locations of holes in the search space. The probability of having particles in these regions is close to zero depending on the importance density.

With these observations in mind, consider the following example. Assume that the first k measurements, Y^k, render the particles $\{X_k(i)\}_{i=1}^{N_1}$ the greatest importance weights out of the entire set, $\{X_k(i)\}_{i=1}^{N}$, where $N > N_1$. Execution of a resampling procedure produces a new set, $\{\tilde{X}_k(i)\}_{i=1}^{N}$, in which all offspring belong to $\{X_k(i)\}_{i=1}^{N_1}$. If the regularization intensity is relatively small, the newly constructed regularized set is rather close to $\{\tilde{X}_k(i)\}_{i=1}^{N}$. Obviously, the new particle set covers a reduced portion of the search space. Now, if the measurements Y^k are located at the tail of $p_{\mathcal{Y}^k}$, it is likely that the following measurements will render the particle set small importance weights. Because the regularization intensity is small, all particles of the next steps will stay in the neighborhood of their parent particles, ever trapped in a small probability region.

Based on this example, one can conclude that retaining some of the underweighted particles of the set $\{X_k(i)\}_{i=N_1+1}^{N}$ would be beneficial in getting away from the trap. In the ensuing paragraphs, two techniques based on this insight are presented.

3. Importance Weight Cooling Schedule

A simple idea for retaining underweighted particles is borrowed from the simulated annealing optimization method [14,15]. The concept of simulated annealing is based on the manner in which metals recrystalize in the process of annealing, which can be thought of as an adiabatic approach to the lowest energy state. If the initial temperature of the system is too low or the cooling process is insufficiently slow, the system may be trapped in a local minimum energy state. Analogously to the simulated annealing "cooling process," the PF algorithm's importance weight update equation is modified to allow some lightly weighted particles to survive via slowing down the algorithm's search for the best particle representation of the posterior PDF. Thus, the modified importance weight update rule is

$$w_k(i) = p_{\mathbf{y}_k|\mathbf{q}_k}(Y_k|\boldsymbol{q}_k(i))^{1-T_k}\tilde{w}_{k-1}(i) \tag{15}$$

where $T_k \in [0, 1]$ is the temperature parameter. When $T_1 = 1$, the particles are considered equally weighted regardless of the incoming observation. Decreasing the temperature allows less underweighted particles to survive. In this work, the following cooling schedule is adopted:

$$T_k = 1 - \gamma(k-1), \qquad k \in \left[1, \text{int}\left(\frac{1}{\gamma}\right) + 1\right] \tag{16}$$

where $\gamma \ll 1$ is some predetermined constant.

4. Partitioned Particle Filtering

Another method for allowing lightly weighted particles to survive is based on partitioning the entire particle set into several smaller sets and using each set as the particle population of a different PF. Regarding the particle's weight as a measure of effectiveness in the PDF approximation while recalling that it is normalized by the sum of all weights, one can conclude that lightly weighted particles can survive by having a PF implemented exclusively using lightly weighted samples. However, implementing a PF using highly unlikely samples may cause a severe numerical problem in the normalization of the particles' weights, when $\sum w_k(i)$ approaches the machine's underflow level, ϵ. For this reason, it is necessary to bound the smallest allowed value for the importance weight, such that

$$w_k(i) \geq w_{min} \geq \epsilon/N, \quad \forall\, i, k \tag{17}$$

Let $S = \{[\boldsymbol{q}_k(j)^T, \boldsymbol{\Omega}_k(j)^T]^T\}_{j=1}^{N}$ be the entire particle set, and consider a partition of this set into m disjoint subsets $\{S_i\}_{i=1}^{m}$, such that $S = \bigcup_{i=1}^{m} S_i$, as shown in Fig. 2. Let $[\hat{\boldsymbol{q}}_k(i)^T, \hat{\boldsymbol{\Omega}}_k(i)^T]^T$ be the minimum mean squared error (MMSE) estimate of the ith particle
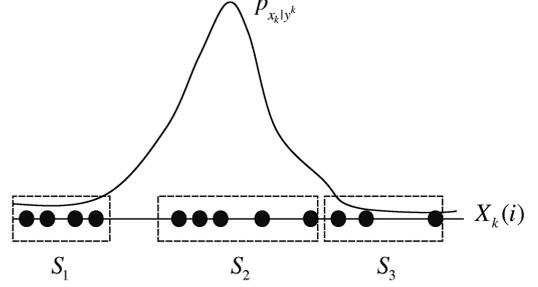

Fig. 2 Illustration of particle set partitioning. Particles and subsets are represented by the bold circles and dashed boxes, respectively. The entire particle set is $S = S_1 \bigcup S_2 \bigcup S_3$. The lightly weighted particles constitute the subsets S_1 and S_3.

filter obtained using the set S_i. The interaction between the particle filters is established by taking the most likely MMSE state estimate trajectory, that is,

$$\left[\hat{\boldsymbol{q}}_k^T, \hat{\boldsymbol{\Omega}}_k^T\right]^T = [\hat{\boldsymbol{q}}_k(i^*)^T, \hat{\boldsymbol{\Omega}}_k(i^*)^T]^T \tag{18a}$$

where

$$i^* = \arg \max_{1 \leq i \leq m} p_{\mathcal{Y}^k|x^k}(Y^k|\hat{X}^k(i)) \tag{18b}$$

and

$$\hat{X}^k(i) \triangleq \{[\hat{\boldsymbol{q}}_1(i)^T, \hat{\boldsymbol{\Omega}}_1(i)^T]^T, \ldots, [\hat{\boldsymbol{q}}_k(i)^T, \hat{\boldsymbol{\Omega}}_k(i)^T]^T\} \tag{19}$$

An equivalent recursive version of Eq. (18b) is written as

$$i^* = \arg \max_{1 \leq i \leq m} W_k(i) \tag{20}$$

where

$$W_k(i) = c_k p_{\mathbf{y}_k|\mathbf{q}_k,\boldsymbol{\omega}_k}(Y_k|[\hat{\boldsymbol{q}}_k(i)^T, \hat{\boldsymbol{\Omega}}_k(i)^T]^T)W_{k-1}(i)$$
$$= c_k p_{\mathbf{y}_k|\mathbf{q}_k}(Y_k|\hat{\boldsymbol{q}}_k(i))W_{k-1}(i) \tag{21a}$$

$$W_1(i) = c_1 p_{\mathbf{y}_1|\mathbf{q}_1,\boldsymbol{\omega}_1}(Y_1|[\hat{\boldsymbol{q}}_1(i)^T, \hat{\boldsymbol{\Omega}}_1(i)^T]^T) = c_1 p_{\mathbf{y}_1|\mathbf{q}_1}(Y_1|\hat{\boldsymbol{q}}_1(i)) \tag{21b}$$

with the normalization constant c_k set to

$$c_k \triangleq \left[\sum_{j=1}^{m} p_{\mathbf{y}_k|\mathbf{q}_k}(Y_k|\hat{\boldsymbol{q}}_k(j))W_{k-1}(j)\right]^{-1} \tag{22}$$

and $W_0(j) \triangleq 1$, $j = 1, \ldots, m$.

5. Filter Reinitialization

Even after implementing the aforementioned techniques, the PF algorithm may still be trapped in a low-probability region (hole) of the search space. In such cases, the filters' unnormalized cost functions $c_k^{-1}W_k(i)$ drop below some threshold value ϵ_1/m. The best solution in this situation is to reinitialize the algorithm. Thus, reinitialization is carried out whenever

$$\sum_{i=1}^{m} c_k^{-1}W_k(i) < \epsilon_1, \quad \forall\, k > k_1 \tag{23}$$

D. Inertia Tensor Estimation via State Augmentation

The attitude and angular-rate PF turns out to be highly sensitive to inertia modeling imperfections. This drawback originates in the angular-rate dynamics characterized by the Euler equation. Experience shows that even a minor inertia deviation on the order of 0.1% in one of the major axes may lead to filter divergence. Therefore, unless the inertia is known to within a sufficient level of

accuracy, proper implementation of the PF should include an appropriate inertia estimation scheme.

A straightforward solution of this problem consists of augmenting the original state by the vectorized form of the spacecraft's inertia tensor. However, the implementation of the PF algorithm for the augmented system is subject to the aforementioned curse of dimensionality. Practically, this implies that any increase in state dimension is associated with a significant increase in the number of particles needed for acceptable PF performance. As a consequence, the PF algorithm may become prohibitively computationally intensive.

In the present case, the dimension of the state vector accounting for inertia imperfections lies between 8 and 13, corresponding to 1–6 inertia entries. In this work, a PF was implemented assuming a 4% error in just one of the diagonal inertia entries. Unfortunately, satisfactory performance could only be obtained using no less than 20,000 particles. The related results are summarized in Sec. IV.

E. On the Viability of the Interlacing Technique

The dimensionality problem is alleviated in the related works of [10,11] via interlacing the PF algorithm with a parameter estimator. Using this approach, the PF maintains its original state dimension, whereas the additional parameters are handled by an external estimator. The information sharing between the two estimators is established in an interlacing manner, that is, each estimator is driven by its companion algorithm's previous-time estimates.

Although the interlacing approach degrades the filtering performance, it is viable and much more computationally efficient than the conventional state-augmentation technique. However, there are some cases in which this method fails to work. The weakness of PF interlacing lies in the underlying assumption of high estimation accuracy gained within a relatively short time. This implies that sufficient information is within reach even though the dependency of the state and parameter channels is misrepresented (the interlacing scheme handles these two channels as independent of each other).

It turns out that, in the case under consideration, the interlacing technique fails to work. This follows from the fact that both information channels, corresponding to the attitude and the angular rate, are tightly coupled due to the physical nature of the problem. Thus, any separation of the two, which amounts to neglecting their physical coupling, causes wrong processing of the information embedded in the measurements, which brings about filter divergence.

IV. Simulation Study

The fast attitude and angular-rate PF algorithm is applied to a realistic LEO spacecraft. The spacecraft is in a near-circular 90 min (350 km) orbit with an inclination of 35 deg (these orbital parameters correspond to the Tropical Rainfall Measuring Mission [4]). The spacecraft's inertia tensor is $J = \mathrm{diag}\{500, 550, 600\}$ Kg · m^2. In all simulation runs, the spacecraft's initial attitude quaternion is randomly sampled from a uniform distribution over the unit three sphere. The initial angular rate is sampled from a uniform distribution, with the norm not exceeding 30 deg/s. In all simulations, the spacecraft's angular rates are numerically integrated using the Dormand–Prince explicit Runge–Kutta (4, 5) formula, implemented in MATLAB®'s ODE45 routine. The spacecraft is equipped with a three-axis magnetometer (TAM) that provides the vector measurements. The TAM noise is modeled as a zero-mean Gaussian white process with a standard deviation of 50 nT. The Earth's magnetic field is modeled using the eighth-order international geomagnetic reference field.

The fast attitude and angular-rate PF is initialized using the OPF of [11]. The OPF is executed during the first 3 min, processing the first batch of measurements. After 3 min, $N_\omega = 30$ angular-rate particles are taken out of the OPF population and used in composing the initial population of the combined attitude and angular-rate PF, as described in Sec. IV.C. The initial quaternion population is produced using the initialization algorithm detailed in [10] with $N_q = 60$

quaternion particles. This yields a total population of $N = N_\omega N_q = 1800$ particles, denoted as $\{[\boldsymbol{q}_0(i)^T, \boldsymbol{\Omega}_0(j)^T]^T, \ i = 1, \ldots, N_q, \ j = 1, \ldots, N_\omega\}$.

The PF uses the importance weight cooling schedule and the particle set partitioning methods described previously. The synthetic temperature constant is set to $\gamma = 5 \times 10^{-4}$.

The particle set consists of four partitions, each containing 450 particles, which amounts to four particle filters running simultaneously. The minimal importance weight and the minimal filter cost function threshold values are set to $w_{\min} = 10^{-12}$ and $\epsilon_1 = 10^{-16}$, respectively. The resampling threshold is set to $N_{\mathrm{th}} = \frac{2}{3}N$. This value is chosen based on tuning runs. Decreasing N_{th} may be beneficial because resampling procedures will be executed less frequently, consequently introducing less Monte Carlo variations into the estimates. However, this increases the algorithm's sensitivity to heavy tailed measurement noise PDFs. The numerical integration of the Euler equation in the evolution stage of the PF is performed using the fourth-order Runge–Kutta method with a constant time interval of $\Delta t/5$. The TAM sampling rate is 1 per 5 s. The attitude estimation error is computed as

$$\alpha_k = 2 \arccos(\delta q_{4_k}) \tag{24}$$

(in degrees), where δq_{4_k} is the scalar component of the error quaternion at time k, $\delta \boldsymbol{q}_k$, defined as

$$\delta \boldsymbol{q}_k \triangleq \boldsymbol{q}_k \otimes \hat{\boldsymbol{q}}_k^{-1} \tag{25}$$

The angular-rate estimation error (in degrees per second) is computed as

$$\|\delta \boldsymbol{\Omega}_k\| \triangleq \|\boldsymbol{\Omega}_k - \hat{\boldsymbol{\Omega}}_k\|_2 \tag{26}$$

A. Precise Inertia Modeling

In the following runs, the inertia tensor is assumed to be known. The results of 500 Monte Carlo runs are presented in Figs. 3 and 4. In these figures, the statistical distributions of the estimation errors are illustrated via their percentile curves. The results are presented starting from time = 3 min subsequent to the ending of the initialization OPF run.

Figure 3 shows the statistical distribution of the quaternion estimation error over the period of 3 h. From this figure it can be clearly seen that, in 95% of the runs, the PF attitude errors reach values of less than 0.02 deg after 3 h.

The PF angular-rate-estimation error statistical distribution is presented in Fig. 4. It can be seen that, in 95% of the runs, the
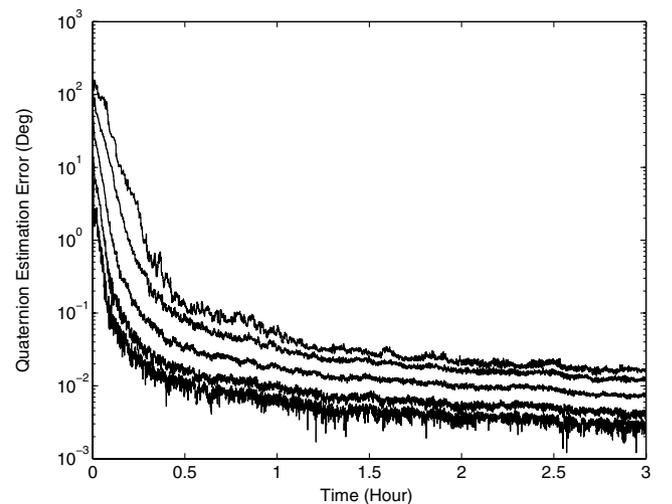

Fig. 3 Statistical distribution of the quaternion estimation error based on 500 Monte Carlo runs. Lines, top to bottom: 95th, 85th, 50th, 15th, and 5th percentiles.
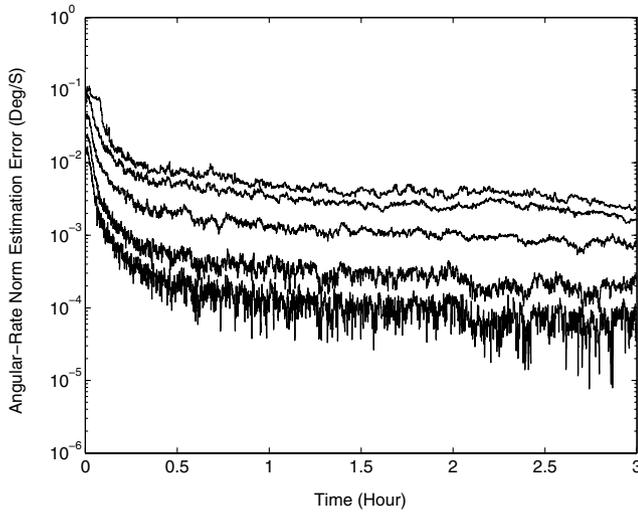
Fig. 4   Statistical distribution of the angular-rate estimation error based on 500 Monte Carlo runs. Lines, top to bottom: 95th, 85th, 50th, 15th, and 5th percentiles.

angular-rate-estimation error reaches values of less than 0.003 deg /s after 3 h.

The partitioned particle filtering selection mechanization is demonstrated using a typical single run in Fig. 5. It can be seen that the filter with the smallest estimation error is selected after about 6 min.

The effect of the partitioned PF scheme is further illustrated in Fig. 6. Shown in this figure is the joint marginal PDF, $p_{q_{1_k}, q_{2_k} | \mathcal{Y}^k}$, after a few iterations. The partitioned sample space consists of four visible peaks corresponding to the particle population of each filter. The true state (marked with an "X" in Fig. 6a) lies within a high-probability region.

The advantage of using the partitioned PF method is manifested in the observed lower reinitialization rate of the partitioned PF vs the unpartitioned algorithm. The unpartitioned PF uses a total of $4 \times 450 = 1800$ particles, and both filters use the importance weight cooling schedule. In 500 Monte Carlo simulation runs, the reinitialization rate observed for the partitioned PF is only 2% vs 17% for the unpartitioned PF (using the same total number of particles).

B.   Inertia Modeling Uncertainties

The attitude and angular-rate PF is modified to cope with inertia uncertainty in just one of its diagonal entries. This is done by taking
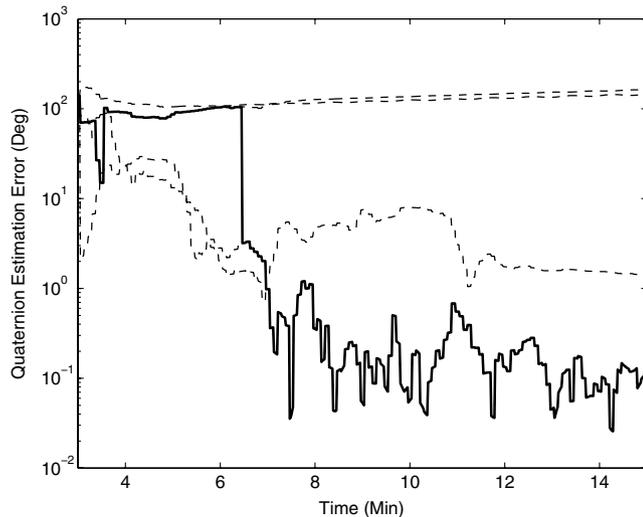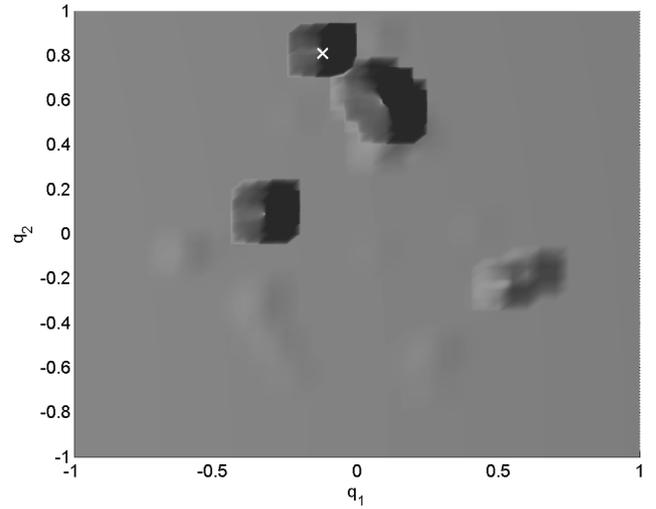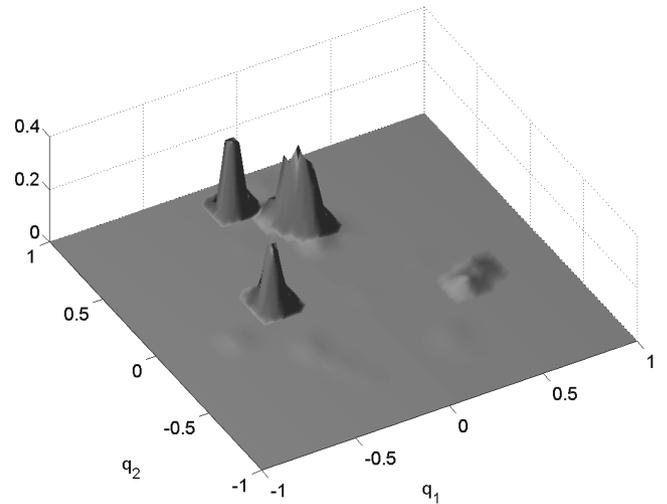

Fig. 5   Partitioned particle filtering selection mechanization (single run): the four particle filters' attitude estimation errors (dashed lines) and the estimation error of the selected filter (bold line).

a) Marginal PDF top view

b) Marginal PDF

Fig. 6   Partitioned PDF of the first two entries of the quaternion, $p_{q_{1_k}, q_{2_k} | \mathcal{Y}^k}$, after a few iterations. True state is marked with an "X."

the first diagonal element of the inertia tensor to be the eighth component of the state vector. In this case, the inertia error is defined as

$$\delta J_k = |J_{11} - \hat{J}_k| \qquad (27)$$

where J_{11} and $\hat{J}_k$ denote the true and the estimated first diagonal inertia element, respectively. The inertia-augmented PF is run using a total number of 20,000 particles. The simulation results, consisting of a single run in which the initial attitude and inertia errors are set to $\alpha_0 = 120$ deg and $\delta J_0 = 20$ Kg · m², are shown in Figs. 7–9. From these figures it can be recognized that the PF converges, almost immediately, to within 0.05 deg in attitude errors. Similarly, the angular-rate norm estimation error attains values of less than 0.02 deg /s. Figure 9 shows that the inertia estimation error reaches values lower than 0.1 Kg · m² after a couple of minutes.

C.   Computational Aspects

To assess the computational requirements of the attitude and angular-rate PF, the filter was run for 2160 cycles (wherein each cycle involves one measurement update and one evolution stage). At the measurement rate of 1 per 5 s, this run corresponds to a real-time duration of 3 h. Implemented in MATLAB® (running in interpreter mode without any code optimization) on a Pentium 4/2.8 GHz machine under the Windows XP operating system, the run required 3.24 h, which means that even this crude implementation of the filter
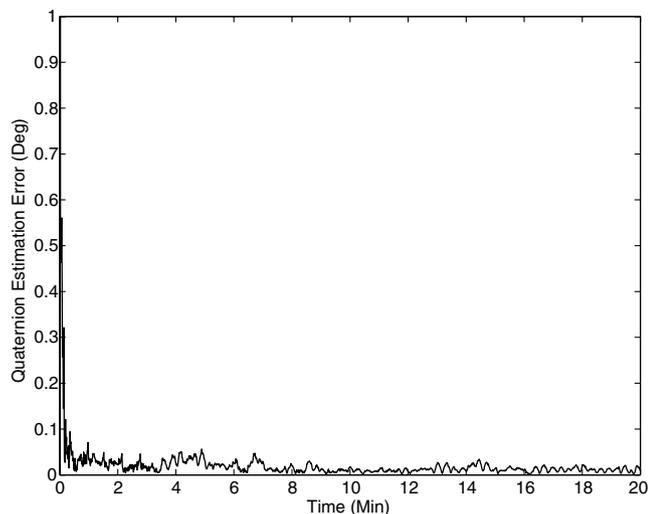
Fig. 7 The quaternion estimation error of the gyroless attitude and inertia PF. Single run.
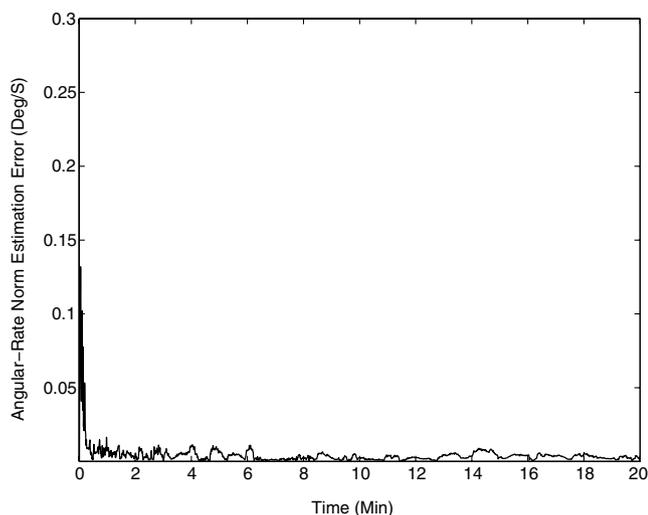

Fig. 8 The angular-rate estimation error of the gyroless attitude and inertia PF. Single run.
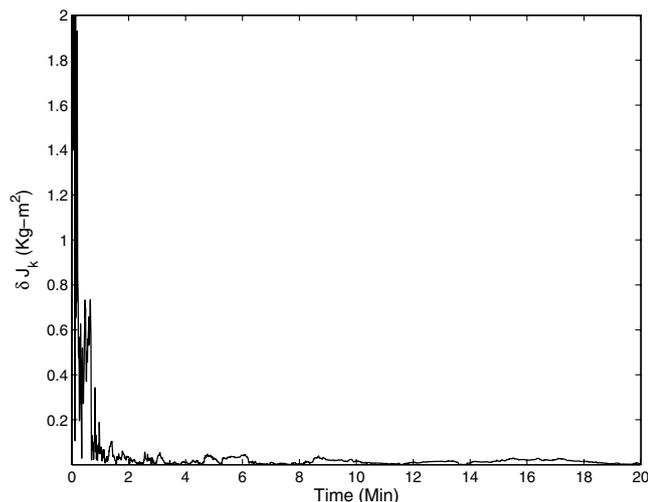

Fig. 9 The inertia estimation error of the gyroless attitude and inertia PF. Single run.

is almost real-time compliant. Clearly, a real-life implementation of the filter (properly coded and compiled in a high-speed computer language under a real-time operating system) should easily close the aforementioned computational requirement "gap," thus permitting the use of the filter in real-time applications. Moreover, recent advances in the parallel implementation of particle filters on field-programmable gate arrays [16,17] should facilitate the use of the newly proposed filter even further.

V. Conclusions

A new particle filtering algorithm is presented for the simultaneous estimation of both the attitude and angular rate from vector observations in gyroless applications. The new estimator naturally copes with non-Gaussian driving processes and with the inherent nonlinearity of the attitude and angular-rate-estimation problem. The dimensionality problem is alleviated by using an efficient initialization procedure, along with an importance weight cooling schedule and particle set partitioning. A simulation study, involving a realistic LEO spacecraft model, is presented, in which the new algorithm demonstrates a fast convergence rate and high estimation accuracy. The new gyroless PF can be modified to cope with inertia tensor uncertainties at the expense of computational efficiency by using the common state-augmentation technique.

References

[1] Wahba, G., "A Least-Squares Estimate of Satellite Attitude. Problem 65-1," *SIAM Review*, Vol. 7, No. 3, July 1965, p. 409.
doi:10.1137/1007077

[2] 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.
doi:10.2514/3.56190

[3] Bar-Itzhack, I. Y., and Oshman, Y., "Attitude Determination from Vector Observations: Quaternion Estimation," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-21, No. 1, Jan. 1985, pp. 128–136.
doi:10.1109/TAES.1985.310546

[4] Crassidis, J. L., and Markley, F. L., "Unscented Filtering for Spacecraft Attitude Estimation," *Journal of Guidance, Control, and Dynamics*, Vol. 26, No. 4, Aug. 2003, pp. 536–542.
doi:10.2514/2.5102

[5] 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.
doi:10.2514/3.56393

[6] Challa, M. S., Natanson, G. A., Deutschmann, J. K., and Galal, K., "A PC-Based Magnetometer-Only Attitude and Rate Determination System for Gyroless Spacecraft," *Proceedings of Flight Mechanics/ Estimation Theory Symposium*, CP 3299, NASA Goddard Space Flight Center, Greenbelt, MD, May 1995, pp. 83–96.

[7] Challa, M., Natanson, G., and Wheeler, C., "Simultaneous Determination of Spacecraft Attitude and Rates Using Only a Magnetometer," *Proceedings of the AIAA/AAS Astrodynamics Specialist Conference*, AIAA, Reston, VA, July 1996.

[8] Natanson, G. A., Challa, M. S., Deutschmann, J. K., and Baker, D. F., "Magnetometer-Only Attitude and Rate Determination for a Gyroless Spacecraft," *Proceedings of The Third International Symposium on Space Mission Operations and Ground Data Systems*, CP 3281, NASA Goddard Space Flight Center, Greenbelt, MD, Nov. 1994, pp. 791–798.

[9] Bar-Itzhack, I. Y., Harman, R. R., and Choukroun, D., "State-Dependent Pseudo Linear Filters for Spacecraft Attitude and Rate Estimation," AIAA, Paper 2002-4461, Aug. 2002.

[10] Oshman, Y., and Carmi, A., "Attitude Estimation from Vector Observations Using a Genetic Algorithm-Embedded Quaternion Particle Filter," *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 4, July–Aug. 2006, pp. 879–891.
doi:10.2514/1.17951

[11] Carmi, A., and Oshman, Y., "Robust Spacecraft Angular Rate Estimation from Vector Observations Using Interlaced Particle Filtering," *Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 6, Nov.–Dec. 2007, pp. 1729–1741.
doi:10.2514/1.28932

[12] Doucet, A., De Freitas, N., and Gordon, N. (eds.), *Sequential Monte Carlo Methods in Practice*, Statistics for Engineering and Information Science, Springer, New York, 2001.

[13] Tortora, P., Oshman, Y., and Santoni, F., "Spacecraft Angular Rate Estimation from Magnetometer Data Only Using an Analytic Solution of Euler's Equations," *Journal of Guidance, Control, and Dynamics*, Vol. 27, No. 3, 2004, pp. 365–373.
doi:10.2514/1.10332

[14] Kirkpatrick, S., Gelatt, C. D., Jr., and Vecchi, M. P., "Optimization by Simulated Annealing," *Science*, Vol. 220, No. 4598, 1983, pp. 671–680.

doi:10.1126/science.220.4598.671

[15] Cerny, V., "Thermodynamical Approach to the Traveling Salesman Problem: An Efficient Simulation Algorithm," *Journal of Optimization Theory and Applications*, Vol. 45, No. 1, 1985, pp. 41–51.
doi:10.1007/BF00940812

[16] Athalye, A., Bolić, M., Hong, S., and Djurić, P. M., "Generic Hardware Architectures for Sampling and Resampling in Particle Filters," *EURASIP Journal on Applied Signal Processing*, Vol. 2005, No. 17, 2005, pp. 2888–2902.
doi:10.1155/ASP.2005.2888

[17] Bolić, M., Djurić, P. M., and Hong, S., "Resampling Algorithms and Architectures for Distributed Particle Filters," *IEEE Transactions on Signal Processing*, Vol. 53, No. 7, July 2005, pp. 2442–2450.
doi:10.1109/TSP.2005.849185