UNDERSTANDING SATELLITE CHARACTERIZATION KNOWLEDGE GAINED FROM RADIOMETRIC DATA

Andrew Harms

Princeton University Kris Hamada, Charles J. Wetterer Pacific Defense Solutions Kim Luu, Chris Sabol Air Force Research Laboratory Kyle T. Alfriend Texas A&M University

Abstract

This paper presents a framework for determining satellite characterization knowledge, in the form of estimated parameter uncertainties, from radiometric observation type, quantity, quality, and in combinations. The approach combines complex forward modeling capability with an Unscented Kalman Filter (UKF) to map observation uncertainties into satellite characterization parameter space. These parameters can include size, shape, orientation, material properties, etc., and the observations can include broadband or narrowband spectral radiometry, spatially resolved or non-resolved imagery, and passive or active optical data. In order to demonstrate the effectiveness of the technique, the example of using photometric light curve observations to estimate the orientation of a cube is presented. This example is chosen since the orientation uncertainty can be analytically traced from basic radiometry equations and compared to the results of the UKF. The uncertainties can also be tested through Monte Carlo analysis in which simulations are performed 10 times in order to compare observed estimation error sample statistics to the uncertainty predicted by the UKF. There are many optical sensors available and proposed to provide satellite characterization information. Understanding the information content in these data, which this approach provides, allows users to predict the amount and type of data required to obtain desired satellite characterization knowledge as well as provides direction for high pay-off future sensor development efforts.

1 Introduction

A need exists to characterize satellites under a variety of conditions using radiometric observations. The types of radiometric observations can include spatially resolved or non-resolved imagery, passive or active observations, broadband or narrowband spectrum, and any of these in combination, among others. The characterization of satellites includes estimating quantities such as size, shape, orientation, and surface material properties along with a measure of the uncertainty in the estimates. The approach we take to this problem is a model-based estimation algorithm. The general idea is to start with some estimator; examples include Kalman filters, particle filters, and batch processors. The estimator combines existing modeling capabilities with the radiometric observations to produce an estimate of the desired satellite parameters. The uncertainties inherent in the observation data and modeling process are also mapped into the parameter space in the form of state estimate uncertainty. Reductions in the parameter estimate uncertainty then provide a measure of information gain from that particular type, quantity, quality, and/or combination of observation data. This approach enables a valuable predictive capability in which specific sensors and observation times may be selected based upon anticipated information gain and to meet specific accuracy requirements.

Several estimators are available and may be suitable for the problem at hand, each providing a different tradeoff between modeling capability and complexity. The estimator of choice in this paper is the Unscented Kalman Filter (UKF) [3]. The estimator includes a satellite model, forward model for the satellite dynamics, and an observation model for the radiometric data. The complexity of each of these can be varied depending on available resources and

required accuracy of the estimates. Part of the problem is determining how complex or simple each of these pieces should be to attain a desired accuracy in the estimate.

Pose estimation is the focused application in this study. Using *a priori* satellite information encapsulated in the satellite model and knowledge of light propagation and scattering physics captured in the observation model, the time-resolved pose of a satellite can be estimated autonomously through each pass from non-resolved radiometry. The benefits of wavelength diversity, as represented by taking observations with different photometric filters, for this application are examined.

2 System Model

The goal of this framework is to accurately characterize a satellite using limited radiometric data. Further, we want to quantify the information gained from particular sources of data. In this paper, the type of observations is spatially non-resolved passive radiometric data, i.e. the sunlight reflected from an orbiting satellite as it passes through our field of view over the earth. We use this data in an UKF that uses a model of the satellite, complex forward modeling of the satellite dynamics, and complex modeling of the optical system in the observation sensor to provide an estimate of the state (or parameters) of the satellite. The possible parameters for estimation are quite varied, but in this paper we focus on the attitude of the satellite and estimate the roll, pitch, and yaw. The coordinate system defined for these states is very important and is discussed below.

2.1 Satellite Model

The model for the satellite used in this paper is a cube with a different material on each surface. This model provides a good balance in complexity as it is simple enough for some amount of verification by physical analysis while the different materials on each face allow us to examine the effects of different materials on the lightcurve. A model of the satellite is shown in Fig. 1. The materials covering the six faces include 1) black paint, 2) white paint, 3) blue paint, 4) mylar, 5) aluminum, and 6) solar cells. In the coordinates described below, the aluminum face is earth-pointing.



Figure 1: Model of the satellite used in this paper. We use a cube-shaped satellite with a different material on each face.

2.2 Dynamics Model

Dynamics in the UKF are controlled by simple Euler angle attitude definitions. No torques or moments of inertia are used to influence attitude dynamics in this process, and we assume the satellite can achieve both the set attitude and the necessary maneuver to change its orientation from one time-step to the next. Even at this basic level, attitudes and base reference frames must be carefully defined and the propagation of Euler angles from one time-step to the next must be handled within the UKF forward model. The attitude is defined in a series of mappings. First, the coordinate axes must be mapped to the satellite body shape. This is typically defined as a model parameter within the New Solid Model (NSM) file where basic shapes are mapped to a coordinate system to form an overall object shape. Each satellite model therefore has a unique mapping of shape to coordinates; the Cube Satellite system is shown in Fig. 1. The second mapping defines the basic attitude profile. Here the object coordinate system is mapped to vectors in a positional system such as the Earth-Centered Inertial (ECI) system. Definitions in this mapping utilize a two vector system; the first vector represents a strict alignment while the second vector represents a constraint condition. An example of this type of mapping is the (Body $z / Body x) \rightarrow (Nadir / Velocity)$, shown graphically in Fig. 2.



Figure 2: Attitude coordinate transformation from a satellite-centered system to the Earth-Centered Inertial system.

Note that the nadir vector and the velocity vector may not be orthogonal. The nadir condition is a strict alignment, while the velocity condition is a constraint so that the Body x lies as much along the velocity vector as possible. In the above case this is handled with the following procedure:

1. The nadir vector, \hat{n} , is defined by the position of the object in ECI (\hat{r} is the vector from the center of the earth to the satellite)

 $\hat{n} = -\hat{r}$

2. The normal to the orbit plane, \hat{o}_n , is defined by taking the cross product of the nadir vector and the exact velocity vector, \hat{v} ,

 $\hat{o}_n = \hat{n} \times \hat{v}$

3. The tangent velocity vector, \hat{t} , is defined by taking the cross product of the orbit normal with the nadir vector

$$\hat{t} = \hat{o}_n \times \hat{n}$$

With these three vectors (nadir, orbit normal, and tangent) the attitude matrix is formed

$$\mathbf{A} = \begin{bmatrix} t_x & t_y & t_z \\ o_{n_x} & o_{n_y} & o_{n_z} \\ n_x & n_y & n_z \end{bmatrix}$$

For this particular attitude, this completes the attitude profile definition and establishes the reference frame within which the UKF estimates the attitude parameters that define object pose as a function of time.

The final step determines the offset from the attitude profile to the final attitude state and is the result of the matrix multiplication of the attitude matrix with the rotation matrix defined by the estimation parameters

$$A_{final} = AE$$

The final attitude matrix is factorized to Euler angles which are input into the observation model. In the "321" sequence the *z*-axis is rotated first, followed by the *y*-axis, and finally the *x*-axis. The Euler angles are thus

$$\begin{aligned} \theta_1 &= atan2(-A_{final_{3,2}}, A_{final_{3,3}})\\ \theta_2 &= asin(A_{final_{3,1}})\\ \theta_3 &= atan2(-A_{final_{2,1}}, A_{final_{1,1}}) \end{aligned}$$

However, in cases where $\theta_2 \rightarrow 90^{\circ}$ there exists a singularity where the first and third rotation end up being about the same direction in space. Under these conditions, the UKF switches to a "313" sequence, and the equations are

$$\begin{aligned} \theta_1 &= atan2(A_{final_{3,1}}, -A_{final_{3,2}}) \\ \theta_2 &= asin(A_{final_{3,3}}) \\ \theta_3 &= atan2(A_{final_{1,3}}, A_{final_{2,3}}) \end{aligned}$$

2.3 Observation Model

The UKF observation model uses the Time domain Analysis Simulation for Advanced Tracking (TASAT) software to provide high-fidelity satellite brightness estimates. TASAT implements ray-tracing algorithms that are capable of determining shadowing and obscuration effects for all components in complex satellite shapes. In addition TASAT is set up to correctly model satellite positions and their relative geometric relationships with the sun and both ground and space-based observers. Matlab-based code is in place as a supplement to the highly validated and documented capabilities found in TASAT. For the UKF implementation, Matlab takes the attitude states for each sigma point calculated by the UKF and writes it to a TASAT input file along with position and velocity information calculated by an external propagator. Each sigma point is then distributed on a parallel processing platform for processing in TASAT. After the parallel simulations have completed, Matlab collects the results and returns the brightness values to the UKF observation model returns pristine exo-atmospheric satellite brightness and is capable of using several established unit systems including astronomical magnitudes or optical cross section (OCS) in the standard Johnson filter bands. This enables the UKF to be easily compared with photometric data collected from telescope systems and calibrated with their standard reduction processes.

3 Unscented Kalman Filter

The estimation algorithm used in this study is the Unscented Kalman Filter (UKF) first proposed by [3]. The UKF is a non-linear extension to the Kalman Filter (KF). Unlike the Extended Kalman Filter (EKF), an alternative non-linear extension to the KF, it does not approximate the non-linearities in the forward model or observation model but instead relies on sampling the Gaussian distribution. It exploits the proposition that it is easier to approximate a Gaussian distribution than it is to approximate a non-linear function. Using the Unscented Transform, the UKF samples the underlying Gaussian distribution and propagates these samples through the full non-linear model in the predict step. The update step is then nearly identical to the standard KF. Previous work has shown that the UKF produces a better estimate of the moments of the distribution up to the fourth order when compared to the EKF for non-linear models [2, 3] and is especially appropriate for the highly non-linear nature of the radiometric model [4].

3.1 Non-linear System Model and Unscented Filtering

The underlying system model utilized by the UKF is described in [1] and summarized here:

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, k) + \mathbf{w}_k$$
$$\mathbf{\tilde{y}}_k = \mathbf{h}(\mathbf{x}_k, k) + \mathbf{v}_k$$

where \mathbf{x}_k is the $n \times 1$ state vector, $\mathbf{\tilde{y}}_k$ is the $m \times 1$ measurement vector, $\mathbf{w}_k \sim \mathcal{N}(0, \mathbf{Q}_k)$ is the process noise, and $\mathbf{v}_k \sim \mathcal{N}(0, \mathbf{R}_k)$ is the measurement noise at the k^{th} time step. Both noise processes are assumed to be white, zeromean Gaussian processes independent of each other and all other variables in the model. The function $f(\cdot)$ represents the model for the system dynamics, and the function $h(\cdot)$ represents the model for the observations. These models are general and can be utilized by other estimation techniques, such as a particle filter. In general for satellite characterization from radiometric data, the functions $f(\cdot)$ and $h(\cdot)$ are highly non-linear. Rather than finding a linear approximation to these functions, the UKF calculates 2n + 1 sigma points around the estimated mean and propagates each of these through the full (non-linear) system model. The action of the algorithm can be broken down into two steps: the predict and the update.

3.2 UKF Predict

The UKF predict step is formulated as follows. At time step k + 1, the sigma points χ_k^i for i = 0, ..., 2n are calculated from the previous state estimate \hat{x}_k^+ and estimated covariance P_k^+ :

$$\chi_{k}^{i} = \begin{cases} \hat{\mathbf{x}}_{k}^{+} & i = 0\\ \hat{\mathbf{x}}_{k}^{+} + \mathbf{\sigma}_{k}^{i} & i = 1, ..., n\\ \hat{\mathbf{x}}_{k}^{+} - \mathbf{\sigma}_{k}^{i-n} & i = n+1, ..., 2n \end{cases}$$

where σ_k^i is the *i*th column of the matrix $\sigma_k = \sqrt{(n+\lambda)P_k^+}$. The parameter λ is a tuning parameter. The matrix \sqrt{P} is the matrix $M = \sqrt{P}$ that satisfies $MM^T = P$ and can be calculated using the Cholesky decomposition. The sigma

points are then propagated through the forward model

$$\chi_{k+1}^i = f(\chi_k^i)$$

and the predicted mean and predicted covariance are calculated

$$\begin{split} \hat{\mathbf{x}}_{k+1}^{-} &= \frac{1}{n+\lambda} \left\{ \lambda \chi_{k+1}^{0} + \frac{1}{2} \sum_{n=1}^{2n} \chi_{k+1}^{i} \right\} \\ \mathbf{P}_{k+1}^{-} &= \frac{1}{n+\lambda} \left\{ \lambda [\chi_{k+1}^{0} - \hat{\mathbf{x}}_{k+1}^{-}] [\chi_{k+1}^{0} - \hat{\mathbf{x}}_{k+1}^{-}]^{T} + \frac{1}{2} \sum_{i=1}^{2n} [\chi_{k+1}^{i} - \hat{\mathbf{x}}_{k+1}^{-}] [\chi_{k+1}^{i} - \hat{\mathbf{x}}_{k+1}^{-}]^{T} \right\} + \mathbf{Q}_{k} \end{split}$$

The sigma points χ_{k+1}^i are then recalulated using the predicted covariance to account for process noise. The *predicted observation* for each sigma point is calculated using the observation model

$$\gamma_{k+1}^i = \mathbf{h}(\boldsymbol{\chi}_{k+1}^i)$$

The mean observation is

$$\hat{\mathbf{y}}_{k+1} = \frac{1}{n+\lambda} \left\{ \lambda \gamma_{k+1}^0 + \frac{1}{2} \sum_{i=1}^{2n} \gamma_{k+1}^i \right\}$$

and the output covariance is

$$\mathbf{P}_{k+1}^{yy} = \frac{1}{n+\lambda} \left\{ \lambda [\mathbf{\gamma}_{k+1}^0 - \hat{\mathbf{y}}_{k+1}] [\mathbf{\gamma}_{k+1}^0 - \hat{\mathbf{y}}_{k+1}]^T + \frac{1}{2} \sum_{i=1}^{2n} [\mathbf{\gamma}_{k+1}^i - \hat{\mathbf{y}}_{k+1}] [\mathbf{\gamma}_{k+1}^i - \hat{\mathbf{y}}_{k+1}]^T \right\}$$

Finally, the innovation covariance is calculated to account for measurement noise

$$\mathbf{P}_{k+1}^{vv} = \mathbf{P}_{k+1}^{yy} + \mathbf{R}_k$$

and the cross-correlation matrix is

$$\mathbf{P}_{k+1}^{xy} = \frac{1}{n+\lambda} \left\{ \lambda [\mathbf{\chi}_{k+1}^0 - \hat{\mathbf{x}}_{k+1}^-] [\mathbf{\gamma}_{k+1}^0 - \hat{\mathbf{y}}_{k+1}]^T + \frac{1}{2} \sum_{i=1}^{2n} [\mathbf{\chi}_{k+1}^i - \hat{\mathbf{x}}_{k+1}^-] [\mathbf{\gamma}_{k+1}^i - \hat{\mathbf{y}}_{k+1}]^T \right\}$$

3.3 UKF Update

After calculating the predicted observation, the UKF *update* step then incorporates the most recent observation measurement \tilde{y}_{k+1} . The *innovation* v_k is the difference of the actual observation and the predicted observation

$$v_{k+1} \equiv \tilde{\mathbf{y}}_{k+1} - \hat{\mathbf{y}}_{k+1}$$

and the (optimal) Kalman gain is

$$\mathbf{K}_{k+1} = \mathbf{P}_{k+1}^{xy} (\mathbf{P}_{k+1}^{vv})^{-1}$$

Finally, the estimated state and covariance are updated

$$\hat{\mathbf{x}}_{k+1}^+ = \hat{\mathbf{x}}_{k+1}^- + \mathbf{K}_{k+1} \mathbf{v}_{k+1}
onumber \ \mathbf{P}_{k+1}^+ = \mathbf{P}_{k+1}^- - \mathbf{K}_{k+1} \mathbf{P}_{k+1}^{\mathbf{v}\mathbf{v}} \mathbf{K}_{k+1}^T$$

The updated state and covariance are then used to start the estimation at the next time-step.

The simulations provided in this paper make use of the UKF algorithm outlined above which utilizes the TASAT software as the observation model. The cube satellite model (as described above) is $1m \times 1m \times 1m$ in size. The dynamics of the satellite are such that it is not changing attitude through the pass with all states at 0°. Therefore, the forward model does not play a large role in this situation. The initial value for the estimated states are also set to 0° and the initial uncertainty is set to a standard deviation (1 σ) of 5°. The process noise is assumed to be 0°, and the measurement noise introduced is 0.01 m²/sterradian.



Figure 3: Lightcurve comparison from the V and IR filter bands. The lightcurve from the IR band has larger OCS values than the V band lightcurve, and the two curves also have different shapes. By utilizing both lightcurves, the UKF can produce an improved estimate of the attitude than if only one lightcurve is utilized.

4 Diversity of Observation Data

The performance of the filter depends highly on the information contained in the observations which in turn is dependent on the quantity, quality, and type of observations. The intuition is that if the provided observations contain more information, then the filter will produce a better estimate of the parameters of the satellite. In this paper, the observations used by the estimator are samples from passive lightcurves (time-resolved radiometry). These observations depend on the specific features (components and material) of the satellite as well as the lighting and viewing geometry. The question asked is "Can we quantify information gain with each data type introduced into the filter?" Because fielding and maintaining sensors are generally significant investments, data collection resources are usually limited. The key to characterizing satellites is to carefully choose the data sources for diversity linked to specific information gain. Some example methods of diversifying the data include filtering the light into different wavelength bands (e.g. Johnson filter bands), using sensors from different physical locations, and using sensors with different polarization filters. In this paper, the focus is on filtering the light into different wavebands.

To begin, an example is illustrative to explain why diversity of wavebands might be helpful. Fig. 3 shows a simulated lightcurve for the cube satellite model in both the visible (V) band and the infrared (IR) bands. First, it is evident that the IR band has larger OCS values than the V band. As a result, the signal-to-noise ratio is larger for the IR band lightcurve, especially at the start of the pass. Upon closer inspection, the features of the two curves are also quite different. The IR band curve has large OCS values at the beginning of the pass, has a sharp drop in the middle (around 04:50), and then slowly decreases until the end. The V band curve, on the other hand, starts with small OCS values, has a slow rise to a peak in the middle, and then decreases to the end of the pass. This data indicates that some part or parts of the satellite are very bright in the IR band at the beginning of the pass while faint in the V band over the same time period. The residual of the estimated states (difference from the truth states) and the 3 σ uncertainty bound, obtained from the covariance estimate, from a UKF utilizing the V band data are shown in Fig. 5. In particular, the estimate of the pitch is much better when the UKF utilizes the IR band data because the uncertainty (represented by the 3 σ curve shown in red) is much smaller. This matches with our intuition that different information can be extracted from the IR band lightcurve in comparison to the V band lightcurve. These UKF estimates and uncertainties, and all those that follow, are averaged over 10 trials each of which uses an independent instance of the noise injected into the lightcurve.

In the rest of this section, two situations are examined. In the first, utilizing data simultaneously from sensors sensitive to five bands of light – IR and V as well as ultraviolet (UV), blue (B), and red (R) – is considered and compared to the single band cases considered above. In the second case, the data from a filter wheel cycing through the same five bands and providing an observation from one band at each time step is compared with the first case.



(a) Residual and uncertainty in roll

(b) Residual and uncertainty in pitch

(c) Residual and uncertainty in yaw

Figure 4: Residual and uncertainty for the state estimates using the V band data. The estimator does not have a good estimate of the pitch as the uncertainty was not decreased much through the pass. The V band contained little information about the pitch.



Figure 5: Residual and uncertainty for the state estimates using the IR band data. The estimator improved its estimate of the pitch state especially. The pitch estimate has the largest uncertainty, but the uncertainty is much smaller than if the estimator uses the V band data. The IR band contains much more data about the pitch than the V band data.

4.1 Five Simultaneous Bands

The lightcurves for the B and UV bands are shown in Fig. 6. The lightcurve for the B band is similar to the V band lightcurve, while the UV lightcurve is much fainter than any of the others. The R band lightcurve is not shown here



Figure 6: Lightcurves for the B and UV filters. The UV band is much fainter than any of the others. The R filter is nearly identical to the B filter.



Figure 7: Residual and uncertainty for the state estimates using the five simultaneous filter bands. The uncertainty is smaller for all three states compared to the single band estimators. The combined information from all five filter bands leads to a better estimate than the information in any single filter band.

but is very similar to the B band and V band lightcurves. The lightcurve for the IR band, shown in Fig. 3(b), is much brighter than any of the other bands over the first half of the pass.

Fig. 7 shows the residual of the state estimates throughout the pass along with the 3σ uncertainty bound for a UKF utilizing data from all five bands (UV, B, V, R, and IR). All three states have very small final uncertainty, and the final estimate residual is very small. In particular, the final residual and uncertainty are smaller compared to a UKF which only utilizes data from a single band. To account for the increased amount of data and to offer a fair comparison, the variance of the noise injected into the lightcurves was increased by a factor of five for the data used by the simultaneous band UKF described in this section. The reason for this change is to account for the increase in the number of independent observations. Each observation decreases the variance of the estimate, so an estimator with access to five times as much data would decrease the variance of the estimate by a factor of five. Note that the uncertainty decreases rapidly at the start of the pass and sees only a small reduction over the second half of the pass and decreases as quickly as the best estimate from the single band cases. The filter is producing a very good estimate of the attitude for this particular example.

4.2 Filter Wheel



Figure 8: Lightcurve for the filter wheel. The odd shape results from using a different filter at each time step.

With the filter wheel observations, the measurements are taken from each filter rotating the wheel at each time step. The first observation comes from the UV filter, the next from the B filter, then the V filter, then the R filter, and finally the IR filter. The observations then start over with the UV filter. In this configuration, there is only one camera to field and operate, but the information gain from wavelength diversity should still be apparent. The filter wheel lightcurve is shown in Fig. 8. The shape of the lightcurve appears odd because the observations come from successively different filters. The very small OCS values of the UV filter produce the pronounced dips while the large OCS values of the IR filter toward the beginning of the pass appear as large spikes. The other three bands have similar



(a) Residual and uncertainty in roll

(b) Residual and uncertainty in pitch

(c) Residual and uncertainty in yaw

Figure 9: Residual and uncertainty for the state estimates using the filter wheel sensor. The uncertainty is larger than that from the filter using five simultaneous bands, but only slightly larger. The information gain comes mostly from the filter bands utilized and not from the way these filter bands are utilized.



Figure 10: Comparison of the 3σ uncertainty, as calculated from the covariance, for the four cases considered: single IR band data, single V band data, five-band simultaneous data, and filter wheel data. The gain from using multiple bands is apparent, and the gain from using five bands simultaneously is small compared to the filter wheel.

OCS values throughout the pass, and the IR values are similar to the B, V, and R OCS values over the latter half of the pass.

The residual of the state estimates is shown in Fig. 9 along with the estimated uncertainty. In comparing Fig. 9 to Fig. 7, we see very similar performance. The residuals are comparable for both cases in all three state estimates and the final residuals are nearly identical. The shape of the 3σ error bounds are also similar throughout the pass. The simultaneous band filter has access to five times as many independent observations which reduces the variance of the estimate, but this is accounted for by increasing the noise injected into the lightcurves as described in the previous section.

As a final comparison, Fig. 10 compares the 3σ uncertainty of the states for each case considered in this paper: single IR band data, single V band data, five simultaneous bands of data, and the filter wheel data. The advantage of using data from multiple bands versus a single band is great. The reduction in uncertainty is also not very large when comparing the five simulataneous bands of data against the filter wheel. The advatage seems to be in using data from diverse sources rather than the exact way in which this data is utilized, at least in this particular case.

5 Conclusion

The model-based estimation approach provides a capability to quantify the value of specific data types for satellite characterization requirements. The results shown in this paper demonstrate application to pose estimation with various filter bands in two different collection configurations. The information content of the various filter bands are quantified and compared against the information content of a particular combination of filter bands. No attempt was made to optimize the combination of filter bands due to scope of the study and the likely dependence on the surface materials

of particular satellites. The primary objectives of this study are focused on development of the model-based estimation capability and demonstration on a simple object. While the analysis offered in this paper is somewhat limited - e.g. one satellite pass, starting the estimator at the truth, stationary pose – the potential for further analysis is very apparent.

There are many ways to expand upon this study. Different types of radiometric data from among the possibilities listed earlier may be examined both singly and in combination with filter bands. Non-radiometric data types, such as polarimetry and radar, may also be incorporated under the same approach, given appropriate observation models. As mentioned before, data diversity may also be achieved through simultaneous data collections from multiple, geographically-separated sites. From a sensor management perspective, this is a very relevant issue addressing the value of scheduling multiple sensors on the same satellite weighed against the greater cost of doing so. Limitations in availability of resources dictate that tradeoffs must be examined routinely. The model-based estimation approach provides an essential capability toward the next-generation paradigm of tasking for information versus the current scheme of tasking for data.

References

- [1] John L. Crassidis and F. Landis Markley. Unscented Filtering for Spacecraft Attitude Estimation. In AIAA Guidance, Navigation, and Control Conference and Exhibit, 2003.
- [2] S. J. Julier, J. K. Uhlmann, and H. F. Durrant-Whyte. A New Approach for Filtering Non-linearSystems. In Proceedings of the American Control Conference, 1995.
- [3] Simon J. Julier and Jeffrey K. Uhlmann. A New Extension of the Kalman Filter to Nonlinear Systems. Proceedings SPIE, 3068:182–193, 1997.
- [4] Charles J. Wetterer and Moriba Jah. Attitude Estimation from Light Curves. Journal of Gudance, Control, and Dynamics, 32:1648–1651, 2009.