Incorporating Target Priorities in the Sensor Tasking Reward Function

Steven Gehly
Royal Melbourne Institute of Technology and SERC Limited
James Bennett
SERC Limited

ABSTRACT

Orbital debris tracking poses many challenges, most fundamentally the need to track a large number of objects from a limited number of sensors. The use of information theoretic sensor allocation provides a means to efficiently collect data on the multitarget system. An additional need of the community is the ability to specify target priorities, driven both by user needs and environmental factors such as collision warnings. This research develops a method to incorporate target priorities in the sensor tasking reward function, allowing for several applications in different tasking modes such as catalog maintenance, calibration, and collision monitoring. A set of numerical studies is included to demonstrate the functionality of the method.

1. INTRODUCTION

With the proliferation of objects occupying near-Earth space over the past several decades, the allocation of limited sensor resources to maintain knowledge of the orbital debris environment has become increasingly important. Recent advances in information theoretic sensor tasking allow formulation of sensor allocation schemes in terms of information gain functionals computed from the multiobject state and hypothesized measurements [1, 2]. These reward functions quantify the difference between the prior and posterior probability density functions (PDFs), and therefore account for the reduction in uncertainty for each object as a result of computing a measurement update. However, as currently formulated, they do not account for individual target priorities, which is a key requirement for the space situational awareness (SSA) problem.

The need for improved tracking of high priority targets may be driven by user requirements, or may arise from the state of the debris environment itself. For instance, if a collision is predicted between two objects, it is naturally important to schedule additional observations to refine the estimates and predicted probability of collision. Target priorities also provide a straightforward way to implement different sensor tasking modes. For example, a sensor in calibration mode may be tasked to observe only certain objects while ignoring others. Similarly, sensors dedicated to collecting data for object characterization may have a different set of object priorities from those tasked with general catalog maintenance.

This research employs a sensor allocation scheme driven by an information theoretic reward function [1, 11], augmented by a tactical importance function (TIF) to specify target priorities [3]. The TIF has been developed in conjunction with Finite Set Statistics (FISST)-based multitarget filters, and offers a mathematically sound approach to the sensor tasking problem including individual target priorities.

2. BACKGROUND

2.1. CPHD FILTER

The simplest FISST-derived estimator is the Probability Hypothesis Density (PHD) filter, which predicts and corrects the first moment of the multitarget PDF, known as the intensity function or PHD. As a first moment approximation, the PHD filter truncates information in the PDF related to target number, and is known to suffer from high variability in the estimated number of targets, particularly in cases where detections of objects are missed due to imperfect sensors [4]. To address this issue, the Cardinalized PHD (CPHD) filter propagates and updates a cardinality distribution in addition to the PHD function [5, 6]. The cardinality is a discrete probability distribution over the possible number of targets.

Incorporating the cardinality recursion produces a more computationally complex filter, but also provides a more stable estimate of the number of objects at each time.

The CPHD filter is based fundamentally on the assumption that the multitarget random finite set (RFS) can be approximated as an independent and identically distributed (i.i.d.) cluster process. The CPHD recursion further assumes that targets evolve and generate measurements independently of one another, that the processes of target survival and birth are independent, and that the measurement and clutter RFSs are independent, with clutter approximated as an i.i.d. process. The full CPHD recursion is generally intractable, but may be simplified by approximating the PHD using a Gaussian Mixture Model (GMM) to yield a closed-form solution. Following the development of Vo et al. [5], the PHD at time t_k is approximated as a weighted sum of J_k Gaussian PDFs,

$$\nu_k(\boldsymbol{x}) \approx \sum_{j=1}^{J_k} w_k^{(j)} p_g(\boldsymbol{x}; \boldsymbol{x}_k^{(j)}, P_k^{(j)}) \tag{1}$$

where $w_k^{(j)} > 0$ are the weights and $p_g(\boldsymbol{x}; \boldsymbol{x}_k^{(j)}, P_k^{(j)})$ is a multivariate Gaussian PDF with mean $\boldsymbol{x}_k^{(j)}$ and covariance $P_k^{(j)}$. The prediction step for the PHD and cardinality is given by

$$\nu_{k|k-1}(\boldsymbol{x}) = \gamma_k(\boldsymbol{x}) + p_S \sum_{j=1}^{J_{k-1}} w_{k-1}^{(j)} p_g(\boldsymbol{x}; \boldsymbol{x}_{k|k-1}^{(j)}, P_{k|k-1}^{(j)}) \tag{2}$$

$$p_{k|k-1}(n) = \sum_{j=0}^{n} p_{\Gamma,k}(n-j) \sum_{l=j}^{\infty} C_j^l p_{k-1}(l) p_S^{(j)} (1-p_S)^{(l-j)} \tag{3}$$

where $\gamma_k(\boldsymbol{x})$ represents the PHD of new, or birth, targets and can be approximated by a GMM as in Eq. (1), $p_{\Gamma,k}(\cdot)$ represents the cardinality of target birth and $C_j^l = \frac{l!}{j!(l-j)!}$ is the binomial coefficient. The probability of target survival, p_S, is assumed to be independent of the target state.

The update equations for the cardinality distribution and PHD are coupled,

$$\nu_k(\boldsymbol{x}) = \frac{\langle \Psi_k^1[w_{k|k-1}, Z_k], p_{k|k-1} \rangle}{\langle \Psi_k^0[w_{k|k-1}, Z_k], p_{k|k-1} \rangle} (1-p_D) \nu_{k|k-1}(\boldsymbol{x}) + \sum_{\boldsymbol{z} \in Z_k} \sum_{j=1}^{J_{k|k-1}} w_k^{(j)}(\boldsymbol{z}) p_g(\boldsymbol{x}; \boldsymbol{x}_k^{(j)}, P_k^{(j)}) \tag{4}$$

$$p_k(n) = \frac{\Psi_k^0[w_{k|k-1}, Z_k](n) p_{k|k-1}(n)}{\langle \Psi_k^0[w_{k|k-1}, Z_k], p_{k|k-1} \rangle} \tag{5}$$

where Z_k is the measurement set, the angle bracket $\langle a, b \rangle$ is used to denote the inner product of a and b, and additional terms are given as follows:

$$\Psi_k^u[w, Z](n) = \sum_{j=0}^{\min(m_k, n)} (m_k - j)! p_\kappa(m_k - j) P_{j+u}^n \frac{\langle 1 - p_D, \nu \rangle^{n-(j+u)}}{\langle 1, \nu \rangle^n} \sigma_j(\Lambda_k(w, Z)) \tag{6}$$

$$\Lambda_k(w, Z) = \left\{ \frac{\langle 1, \kappa_k \rangle}{\kappa_k(\boldsymbol{z})} p_D \sum_{j=1}^{J_{k|k-1}} w_{k|k-1}^{(j)} p_g(\boldsymbol{z}; \boldsymbol{z}_k^{(j)}, P_{zz}^{(j)}) : \boldsymbol{z} \in Z \right\} \tag{7}$$

$$w_k^{(j)}(\boldsymbol{z}) = p_D w_{k|k-1}^{(j)} p_g(\boldsymbol{z}; \boldsymbol{z}_k^{(j)}, P_{zz}^{(j)}) \frac{\langle \Psi_k^1[w_{k|k-1}, Z_k \backslash \{\boldsymbol{z}\}], p_{k|k-1} \rangle}{\langle \Psi_k^0[w_{k|k-1}, Z_k], p_{k|k-1} \rangle} \frac{\langle 1, \kappa_k \rangle}{\kappa_k(\boldsymbol{z})} \tag{8}$$

where m_k is the number of measurements, $p_\kappa(\cdot)$ is the cardinality of clutter, $\kappa_k(\cdot)$ is the PHD of clutter, $\sigma_j(\cdot)$ are the elementary symmetric functions, and $\boldsymbol{z}_k^{(j)}$ and $P_{zz}^{(j)}$ are the predicted mean and covariance of component j in the measurement space. The probability of detection, p_D, is assumed to be independent of the target state; however, this assumption will be revisited in the following section.

At time t_k, the number of targets can be estimated using the PHD or the estimated cardinality distribution.

$$N_k = \int \nu_k(\boldsymbol{x}) d\boldsymbol{x} \approx \sum_{j=1}^{J_k} w_k^{(j)} \tag{9}$$

$$N_k = \sum_{n=1}^{\infty} n \cdot p_k(n) \tag{10}$$

The PHD update Eq. (4) contains two terms, the first of which accounts for the possibility of missed detections. In this case, the filter reweights the *a priori* PHD function. In the GMM approximation, the *a priori* component means and covariances are kept and are not updated by any measurement, and the updated component weights are computed from

$$w_k^{(j)} = \frac{\langle \Psi_k^1[w_{k|k-1}, Z_k], p_{k|k-1} \rangle}{\langle \Psi_k^0[w_{k|k-1}, Z_k], p_{k|k-1} \rangle} (1 - p_D) w_{k|k-1}^{(j)} \tag{11}$$

This term produces $J_{k|k-1}$ GMM components whose weights are scaled by a factor including $(1 - p_D)$. The second term of Eq. (4) accounts for detected objects, and computes a measurement update for each GMM component using each measurement, thereby producing $J_{k|k-1} \cdot m_k$ components. The updated weights for these components are scaled by a factor including p_D and the individual likelihood of measurement to component association $p_g(\boldsymbol{z}; \boldsymbol{z}_k^{(j)}, P_{zz}^{(j)})$. The net result of the measurement update step is the creation of $J_{k|k-1} \cdot (m_k + 1)$ components, many of which will have small weights and not contribute significantly to the GMM approximated PHD.

To keep the problem computationally tractable, components are removed or merged based on user-defined thresholds [7]. In this research, a component is removed if its weight is below a threshold T. Subsequently, if the Mahalanobis distance between two components i and j is less than a threshold U, the merged component weight, mean, and covariance are computed from

$$\bar{w} = w^{(i)} + w^{(j)} \tag{12}$$

$$\bar{\boldsymbol{x}} = \frac{1}{\bar{w}} \sum_{l \in \{i,j\}} w^{(l)} \boldsymbol{x}^{(l)} \tag{13}$$

$$\bar{P} = \frac{1}{\bar{w}} \sum_{l \in \{i,j\}} w^{(l)} [P^{(l)} + (\boldsymbol{x}^{(l)} - \bar{\boldsymbol{x}})(\boldsymbol{x}^{(l)} - \bar{\boldsymbol{x}})^T] \tag{14}$$

The individual component means, covariances, and predicted measurements in Eqs. (2)-(8) are computed using the scaled unscented transform [8,9]. In addition to the GMM components of the PHD, the filter maintains a list of object identifiers from the public Two-Line Element (TLE) catalog. Anytime a new component is added to the GMM as a result of the measurement update, the identifier of the original component is added to the list as well. The approach does not ensure a rigorous treatment of unique track labels incorporated in the multitarget state [3], but in the case that no target birth or death is considered, it does provide a practical means to reconstruct the time history of components and compute the average state errors for each object as described in Section 3.

In this paper, the cardinality is initialized using a binomial distribution [4],

$$p_0(n) = C_n^J (1 - q)^J \left(\frac{q}{1 - q} \right)^n \tag{15}$$

where J is the number of terms in the cardinality distribution and q is the individual probability of existence for each object. If the true number of objects is N and the distribution is initialized with the value $q = N/J$, then $\underset{n}{\operatorname{argmax}}\, p_0(n) = N$ and the estimated number of targets computed from Eq. (10) is $N_k = N$.

2.1.1. Probability of Detection Calculation

Computing the correct probability of detection plays an important role in maintaining custody of objects in the filter. The previous section applies a constant p_D for all components, but this is clearly problematic for objects that are outside the sensor field of view (FOV) at a particular measurement time. In the case that objects are known to be outside the FOV and cannot be detected, assigning the same p_D as for objects in the FOV will quickly cause GMM

components to be downweighted and knowledge of the object to be lost. Meaningful values of p_D need to be assigned for each object, properly accounting for their potential positions with respect to the FOV.

A simple approach to determining $p_D^{(j)}$ for an individual component is to use the predicted measurement and an indicator function [10]. The probability of detection can be modeled as a product of two terms, a constant $p_{D,\text{sensor}}$ that accounts for the sensor's imperfect ability to detect objects in the FOV, and a state-dependent term

$$p_{D,\text{FOV}}^{(j)} = \begin{cases} 1 & \text{if } \boldsymbol{z}^{(j)} \in \text{FOV} \\ 0 & \text{if } \boldsymbol{z}^{(j)} \notin \text{FOV} \end{cases} \tag{16}$$

where $\boldsymbol{z}^{(j)}$ is the predicted measurement for component j computed using the unscented transform. Assuming the two processes are independent, the overall $p_D^{(j)} = p_{D,\text{sensor}} \cdot p_{D,\text{FOV}}^{(j)}$.

This approach addresses the issue, but may not be sufficient in cases where the object is near the edge of the field of view. A solution to this issue is to split GMM components near the edge of the FOV into several new components prior to computing the probability of detection [11]. This allows the means of some components to lie inside the FOV and others outside, thereby accounting for the possibility of a missed detection while still maintaining a high p_D in the FOV. A simple test for splitting is to integrate the component PDF and check against a threshold, for example:

$$0.05 \leq \int_{\text{FOV}} p_g(\boldsymbol{z}; \boldsymbol{z}_k^{(j)}, P_{zz}^{(j)}) d\boldsymbol{z} \leq 0.95 \tag{17}$$

where $\boldsymbol{z}_k^{(j)}$ and $P_{zz}^{(j)}$ are the predicted mean and covariance in the measurement space computed using the unscented transform. If the integral is between the inequalities, the component is taken to be partially in the FOV and split according to a pre-determined library [12]. Having determined a value of p_D for each component, minor changes must be applied in the measurement update equations above. In Eqs. (4) and (8), the appropriate $p_D^{(j)}$ must be used for each component, more easily seen in Eq. (11) for the missed detection case. In Eq. (7), the $p_D^{(j)}$ term must be moved inside the summation, and finally the modification to Eq (6) is most easily simplified by writing out the inner product for the GMM approximation.

$$\langle 1 - p_D^{(j)}, \nu \rangle = \sum_{j=1}^{J} (1 - p_D^{(j)}) w^{(j)} \tag{18}$$

2.2. DEVELOPMENT OF INFORMATION GAIN EQUATIONS

2.2.1. Rényi divergence

This section examines the use of Rényi α-divergence as applied to the SSA sensor allocation problem. The general form of the Rényi divergence is given by [1],

$$R(\boldsymbol{u}) = \frac{1}{\alpha - 1} \log \int f_1(X; \boldsymbol{u})^{\alpha} f_0(X)^{1-\alpha} dX \tag{19}$$

where $\boldsymbol{u}$ is the sensor control vector, and $f_0(\cdot)$ and $f_1(\cdot)$ denote the prior and posterior multitarget PDFs, respectively. For the CPHD filter, the GMM approximation of the intensity function yields the following solution [11].

$$\nu_0(\boldsymbol{x}) \approx \sum_{i=1}^{J_0} w_i p_g(\boldsymbol{x}; \boldsymbol{m}_i, P_i) \quad \nu_1(\boldsymbol{x}; \boldsymbol{u}) \approx \sum_{j=1}^{J_1} w_j p_g(\boldsymbol{x}; \boldsymbol{m}_j, P_j) \tag{20}$$

$$R_{\text{C}}(\boldsymbol{u}) \approx -2 \log \sum_{n \geq 0} \left(\frac{p_1(n; \boldsymbol{u})}{N_1^n} \right)^{1/2} \left(\frac{p_0(n)}{N_0^n} \right)^{1/2} \cdot \left[\int \left(\sum_{i=1}^{J_0} \sum_{j=1}^{J_1} w_i w_j K_{i,j} p_g(\boldsymbol{x}; \boldsymbol{m}_{i,j}, P_{i,j}) \right)^{1/2} d\boldsymbol{x} \right]^n \tag{21}$$

$$K_{i,j} = p_g(\boldsymbol{m}_i; \boldsymbol{m}_j, P_i + P_j)$$

$$P_{i,j} = \left[P_i^{-1} + P_j^{-1} \right]^{-1}$$

$$\boldsymbol{m}_{i,j} = P_{i,j} \left[P_i^{-1} \boldsymbol{m}_i + P_j^{-1} \boldsymbol{m}_j \right]$$

where the subscripts 0 and 1 refer to the prior and posterior distributions respectively, and the value $\alpha = 0.5$ has been used as it provides the best discrimination between PDFs [1, 13]. In the case of the PHD filter, the multitarget PDF is modeled as a Poisson process, and the corresponding Rényi divergence can be derived from Eq. (21) by substituting a Poisson cardinality distribution $p(n) = \dfrac{e^{-\lambda}\lambda^n}{n!}$ and noting that in the PHD filter the Poisson mean number of targets $\lambda = N$.

$$R_{\mathrm{P}}(\boldsymbol{u}) \approx \sum_{i=1}^{J_0} w_i + \sum_{j=1}^{J_1} w_j - 2\int \left(\sum_{i=1}^{J_0}\sum_{j=1}^{J_1} w_i w_j K_{i,j} p_g(\boldsymbol{x}; \boldsymbol{m}_{i,j}, P_{i,j}) \right)^{1/2} d\boldsymbol{x} \tag{22}$$

While the final equation is simpler, note that both the PHD and CPHD Gaussian mixture forms of Rényi divergence require numerical integration.

2.2.2. Hedging

The previous section makes use of a simplified notation that requires some further discussion. The information functional $R(\boldsymbol{u})$ is also a function of the prior and posterior multitarget PDFs in the measurement space. When computing the functional, the prior PDF $f_0(X)$ is propagated to the desired time t_k. A candidate posterior PDF is generated from the CPHD update equations using a measurement set Z_k, which is dependent on the sensor control vector $\boldsymbol{u}$. The information functional is therefore denoted as $R(\boldsymbol{u}, f_0(X), Z_k(\boldsymbol{u}))$. The desire is to select a sensor control vector that maximizes the information gain; however, the realized measurement set $Z_k(\boldsymbol{u})$ is unknown until this selection is made. The unknown measurement set can be eliminated from the information functional by taking the expected value with respect to the measurement set $Z_k(\boldsymbol{u})$, a process known as hedging [3].

$$R(\boldsymbol{u}, f_0(X)) = \mathbb{E}[R(\boldsymbol{u}, f_0(X), Z_k(\boldsymbol{u}))] \tag{23}$$

The expected value can be computed using multiple representative samples of measurement sets $Z_k(\boldsymbol{u})$ based on the given clutter intensity, probability of detection, and measurement likelihood. A simpler approach, used in this work and elsewhere in the literature [1, 2, 14], is to compute Eq. (23) using a single sample in which $Z_k(\boldsymbol{u})$ is taken to be the predicted ideal measurement set (PIMS) [15]. In this case, a single measurement set is computed assuming no clutter or measurement noise and applying $p_D = 1$ for all objects in the FOV, as determined from the component means. The PIMS measurement set is therefore

$$Z_k^{\mathrm{PIMS}}(\boldsymbol{u}) = \bigcup_{\boldsymbol{z}_k^{(j)} \in \mathrm{FOV}} \{\boldsymbol{z}_k^{(j)}\} \tag{24}$$

Use of the PIMS measurement set is an approximation to the expected value [3], such that the expected information gain is given by

$$R(\boldsymbol{u}, f_0(X)) = \mathbb{E}[R(\boldsymbol{u}, f_0(X), Z_k(\boldsymbol{u}))] \cong R(\boldsymbol{u}, f_0(X), Z_k^{\mathrm{PIMS}}) \tag{25}$$

The approximation is applied by simulating a measurement at the mean of each of the N_{FOV} highest weighted components. For example, if the FOV contains 10 components representing 3 estimated objects, then 3 measurements are simulated and located at the 3 largest weighted component means. The use of the highest weighted components is consistent with the approach used for state extraction in Section 3 and in the literature [5]. The updated cardinality and PHD are computed using Eqs. (4)-(5), and the Rényi divergence is computed numerically using Monte Carlo integration.

2.2.3. Incorporating Target Priorities

In order to identify targets of importance (TOIs) within the context of FISST, Mahler has defined the tactical importance function, $\tau(\boldsymbol{x}) \in [0, 1]$ [3, 4]. Following the development of the Posterior Expected Number of Targets of Interest (PENTI) tasking reward function [16, 17], this research seeks to incorporate the TIF in the previously defined information gain equations for the PHD and CPHD filters.

Beginning with the PHD function, the TOI-biased PHD and expected number of targets can be defined from Eqs. 25.225-25.226 of Reference 3:

$$\nu_{\mathrm{TOI}}(\boldsymbol{x}) = \tau(\boldsymbol{x}) \cdot \nu(\boldsymbol{x}), \quad N_{\mathrm{TOI}} = \int \tau(\boldsymbol{x}) \cdot \nu(\boldsymbol{x}) d\boldsymbol{x} \tag{26}$$

In terms of the Gaussian mixture approximation, the equations are assumed to take a form similar to the state dependent probability of detection, in which each GMM component has a specified TIF value τ_j. The approximate forms of the TOI-biased PHD and expected number of targets are then given by:

$$\nu_{\text{TOI}}(\boldsymbol{x}) \approx \sum_{j=1}^{J} \tau_j w_j p_g(\boldsymbol{x}; \boldsymbol{m}_j, P_j), \quad N_{\text{TOI}} \approx \sum_{j=1}^{J} \tau_j w_j \tag{27}$$

For the CPHD filter, it is also necessary to form the cardinality distribution, which in this research is assumed to follow a multi-Bernoulli distribution. The standard and TOI-biased cardinality distributions are given by

$$p(n) = \prod_{j=1}^{J}(1-w_j) \cdot \sigma_n\left(\left\{\frac{w_1}{1-w_1}, \ldots, \frac{w_J}{1-w_J}\right\}\right) \tag{28}$$

$$p_{\text{TOI}}(n) = \prod_{j=1}^{J}(1-\tau_j w_j) \cdot \sigma_n\left(\left\{\frac{\tau_1 w_1}{1-\tau_1 w_1}, \ldots, \frac{\tau_J w_J}{1-\tau_J w_J}\right\}\right) \tag{29}$$

where $\sigma_n(\cdot)$ are the elementary symmetric functions, and the maximum value $q_j = 0.999$ is used for any case $q_j \geq 1$ where q_j is the weight or product of the weight and TIF as appropriate.

With these equations, it is possible to substitute in either the PHD or CPHD form of the Rényi divergence to get the TOI-biased version of the information functional. Sensor tasking decisions are therefore a matter of evaluating potential pointing assignments, simulating measurements per the PIMS approximation, and computing the expected measurement update and value of the reward function. Targets of interest can be specified by assigning values $\tau_j \approx 1$ while all non-TOIs are assigned smaller TIF values. This paper maintains values of the TIF for each component through the filter prediction and correction step, i.e., the values are component-specific and do not change over time. As will be demonstrated in subsequent testing, adjusting the relative TIF values of TOIs and non-TOIs drives sensor tasking to collect more measurements and improve state estimates of higher priority targets.

2.2.4. Analysis of Methods

In order to examine the behavior of the proposed information gain functionals, the single target problem is considered, for which an analytic solution of the Rényi divergence is available. Assuming an object represented by a single Gaussian component and noting that under the PIMS approximation no adjustment is made to the mean, $\boldsymbol{m}_0 = \boldsymbol{m}_1$, the preceding divergence functions for TOIs reduce to

$$R_{\text{C,TOI}} \approx -2\log \sum_{n\geq 0}\left(\frac{p_1(n)}{N_1^n}\right)^{1/2}\left(\frac{p_0(n)}{N_0^n}\right)^{1/2}(w_0\tau_0 w_1\tau_1)^{n/2}\left(\frac{|4P_{0,1}|}{|P_0+P_1|}\right)^{n/4} \tag{30}$$

$$R_{\text{P,TOI}} \approx w_0\tau_0 + w_1\tau_1 - 2\sqrt{w_0\tau_0 w_1\tau_1}\left(\frac{|4P_{0,1}|}{|P_0+P_1|}\right)^{1/4} \tag{31}$$

where the subscripts 0 and 1 denote the prior and posterior distributions respectively, and $P_{0,1}$ is the covariance of the product of the PDFs computed from Eq. (21). The cardinality and estimated number of targets are computed from Eqs. (27) and (29).

$$p_k(n) = \begin{bmatrix}(1-\tau_k w_k) & \tau_k w_k\end{bmatrix}, \quad N_k = \tau_k w_k \tag{32}$$

If it is further assumed that no update is made to the weights or TIF such that $w_0 = w_1 = w$ and $\tau_0 = \tau_1 = \tau$, the equations reduce to:

$$R_{\text{C,TOI}} \approx -2\log\left[(1-\tau w) + \tau w\left(\frac{|4P_{0,1}|}{|P_0+P_1|}\right)^{1/4}\right] \tag{33}$$

$$R_{\text{P,TOI}} \approx 2\tau w\left[1 - \left(\frac{|4P_{0,1}|}{|P_0+P_1|}\right)^{1/4}\right] \tag{34}$$

Note that for each formula, the minimum value of zero is found when $P_1 = P_0$. Alternately, the expected information gain is zero if either w or τ is zero. As a final note, if a Poisson cardinality is substituted in place of the multi-Bernoulli distribution, using a mean value $\lambda = \tau w$, the CPHD equation reduces to the PHD form.

3. NUMERICAL SIMULATION

3.1. LINEAR MOTION EXAMPLE

To examine the behavior and relative performance of the proposed information gain functionals, a simple tasking scenario is considered in which the sensor must select one out of two objects to observe at each time, see Fig. 1. The objects are separated such that the initial Mahalanobis distance is 10, therefore only one object is considered to be in the field of view at each time and the single target information gain equations are applied. At each time, the sensor selects one of the two objects to observe and computes a bearing measurement and corresponding covariance update using the Conventional Kalman Filter update [18]. The selection is based on maximizing the expected information gain for the single object observed, as computed using the CPHD form of the TOI-biased Rényi divergence.

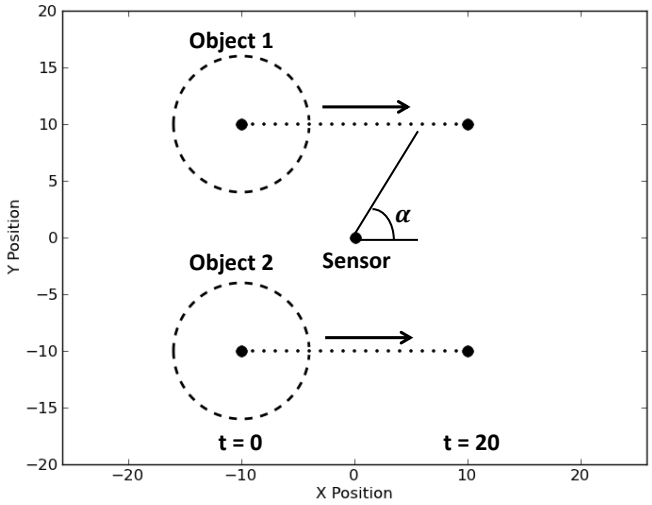

Fig. 1. Linear Motion Test Case

The objects are modeled as moving with constant velocity, with state vector given by $\boldsymbol{x} = [x \quad y \quad \dot{x} \quad \dot{y}]^T$. Over the course of 20 time steps, the objects move past a sensor, which computes bearing measurements $\alpha = \tan^{-1}\left(\frac{y-y_s}{x-x_s}\right)$. The covariance is propagated using a state transition matrix, and the hypothetical measurement update is computed assuming a measurement noise covariance $R_k = 1 \text{ deg}^2$.

Both objects are initialized with the same starting covariance and at all times are the same distance from the sensor, with corresponding equal magnitude measurement α. In the case that both objects have the same priority, the desired sensor management behavior is to alternately select each object, thereby producing an equal number of measurements for each and driving the final uncertainties to similar values. In the case that one object is designated as a TOI, it is expected to have more measurements scheduled and a reduced uncertainty as compared to the non-TOI.

Several cases are considered for different values of the TIF for each object. In all cases, the priority of Object 1 is set as $\tau_1 = 1$. The value for Object 2 is adjusted, from $\tau_2 = 1$ down to $\tau_2 = 0$. The results of the first simulation are provided in Figure 2(a), in which Object 2 is given the same priority as Object 1. The top plot provides the trace of the covariance for each object. Below that are values of the expected information gain, and finally, the identity of the object scheduled for observation at each time. From the plots, it is clear that the sensor alternates between observing the two objects, and the covariance updates follow similar trends, reaching the same value at the final time. On the right, Figure 2(b) provides results for a case where $\tau_2 = 0.5$. While it is clear the expected information gain for Object 2 has been reduced as a result of its lower priority, the difference is not enough to significantly affect tasking. Other than one scan near the end where Object 1 is scheduled twice in a row, the reward function still drives tasking to alternate between the two objects, and produces a similar covariance reduction for both.

Figure 3 provides results for cases in which τ_2 is reduced further to cause the sensor to favor scheduling Object 1. On the left, setting $\tau_2 = 0.1$ is observed to produce many more measurements for Object 1, and a corresponding growth in the covariance for Object 2. Around the halfway point, the covariance for Object 1 has converged to a small

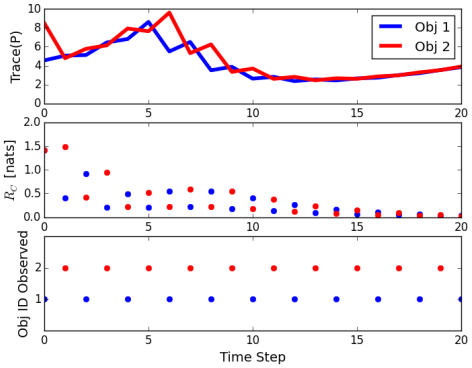

(a) Tasking Results $\tau_2 = 1$

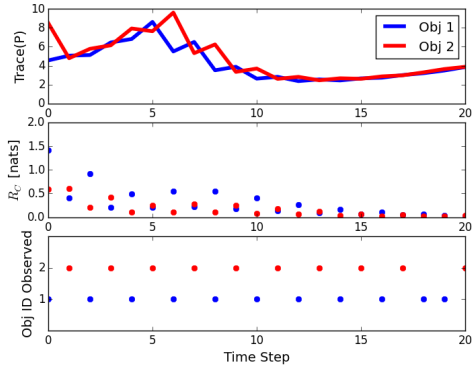

(b) Tasking Results $\tau_2 = 0.5$

Fig. 2. Linear Motion Model Results

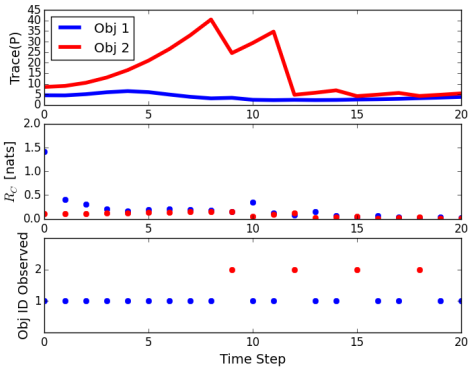

(a) Tasking Results $\tau_2 = 0.1$

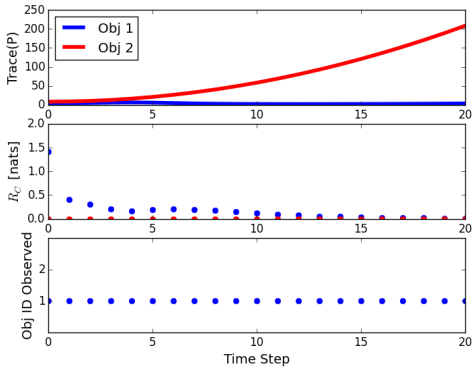

(b) Tasking Results $\tau_2 = 0$

Fig. 3. Linear Motion Model Results

value, and the expected information gain is low, allowing for occasional measurements of Object 2 to be scheduled. Finally, on the right, the priority for Object 2 is set to zero, producing a corresponding value of zero for the expected information gain. All measurements are scheduled for Object 1, with the expected divergence between the covariance values of the two objects.

3.2. ORBITAL DEBRIS SCENARIO

3.2.1. Test Case Description

To verify the performance of the priority driven tasking in an orbital debris scenario, a large scale simulation is included, in which tracking is scheduled for a total of 1118 objects in near-geosynchronous orbits over the course of 24 hours using five GEODSS-like sensors. The objects are selected from the public TLE catalog for the date February 27, 2013, and constrained by the following criteria:

$$0 \leq e \leq 0.3 \qquad 0^\circ \leq i \leq 70^\circ \qquad 0.9 \leq n_m \leq 1.1 \left[\frac{\text{rev}}{\text{sidereal day}}\right]$$

Approximately 10% of the catalog, 112 objects, are randomly selected and designated as TOIs. The test is structured similarly to the linear model, in which the priorities of non-TOIs are reduced to drive additional measurements and improved state estimation for the high priority targets. The same set of TOIs are used in each case.

The ground stations are modeled using the parameters provided in Table 1 [19]. The measurements used are topocentric right ascension and declination,

$$\alpha = \tan^{-1}\left(\frac{y - y_{\text{si}}}{x - x_{\text{si}}}\right) \quad \delta = \sin^{-1}\left(\frac{z - z_{\text{si}}}{\rho}\right) \tag{35}$$

where all values are given in the Earth-Centered Inertial (ECI) frame, $\rho = \sqrt{(x - x_{\text{si}})^2 + (y - y_{\text{si}})^2 + (z - z_{\text{si}})^2}$ is the range, and the si subscript denotes a ground station coordinate. Values for measurement noise and the sensor FOV are provided in Table 2 [20]. The offset in hours from GMT is provided for the date February 27, 2013 used as the initial epoch in the simulation; each sensor is modeled as active for 12 hours per night, from local 6PM to 6AM.

Table 1. GEODSS Sensor Parameters [19]

Site	Latitude [deg]	Longitude [deg]	Altitude [m]	Az Limits [deg]	El Limits [deg]
Socorro, NM	33.82	-106.66	1510.2	[0,360]	[20,90]
Maui, HI	20.71	-156.26	3058.2	[0,360]	[20,90]
Diego Garcia	-7.41	72.45	-61.2	[0,360]	[20,90]
Lisbon, Portugal	38.74	-9.14	77.0	[0,360]	[20,90]
Alice Springs, Australia	-23.70	133.87	583.0	[0,360]	[20,90]

Table 2. GEODSS Sensor Parameters [20]

Site	FOV Size $[\alpha, \delta]$	Noise $[\alpha, \delta]$	GMT Offset [hours]
Socorro, NM	$[1.61^\circ, 1.23^\circ]$	[2″,2″]	-7
Maui, HI	$[1.61^\circ, 1.23^\circ]$	[2″,2″]	-10
Diego Garcia	$[1.61^\circ, 1.23^\circ]$	[2″,2″]	+6
Lisbon, Portugal	$[1.61^\circ, 1.23^\circ]$	[2″,2″]	0
Alice Springs, Australia	$[1.61^\circ, 1.23^\circ]$	[2″,2″]	+9.5

Figure 4 provides a visual representation of each ground station and its full field of regard. The five ground stations are capable of global coverage of the GEO belt with a 20° elevation mask applied. Note that due to their latitude, the Socorro and Lisbon FORs extend over the north pole, and that there are large overlaps in the coverage from all sensors. Object locations at the epoch time are also plotted, from which it is clear that most objects are located near the equator.

Table 3. Initial State Uncertainties

Parameter	Value
σ_a	1.0 km
σ_e	10^{-4}
σ_i	0.01°
σ_Ω	0.01°
σ_ω	0.01°
σ_M	0.01°

The filter is initialized with a PHD function of 1118 components with the initial uncertainty for each specified in orbital elements in Table 3. The mean estimated state of each object is randomly perturbed from the truth using these values, then both the mean and covariance are converted to Cartesian ECI coordinates using the unscented transform [9]. The objects are propagated assuming two-body dynamics, with perturbing forces modeled by a 2x2 spherical harmonics gravity field based on the EGM2008 model [21], solar radiation pressure (SRP), and luni-solar perturbations using the JPL design ephemeris 430 [22,23]. The SRP force is modeled assuming all objects are spherical with area-to-mass ratio 0.05 kg/m^2 and reflectivity $C_r = 1.5$. No unmodeled accelerations are included in the filter. The covariance is assumed Gaussian and propagated using the unscented transform. A simple process noise model is used, in which a diagonal process noise covariance matrix $Q = Q_0 \Delta t$ is added to the predicted covariance at each

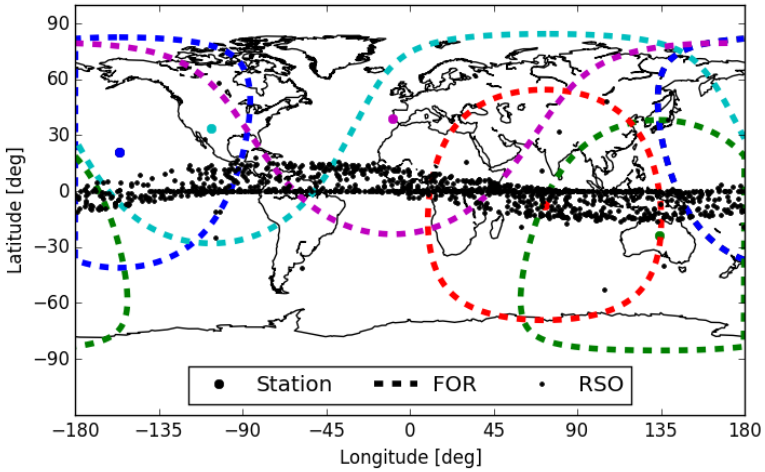

Fig. 4. Ground Station Coverage and Object Locations at Epoch

time. The matrix Q_0 uses standard deviations of 10^{-4} km and 10^{-7} km/s for each of the positions and velocities in ECI, and Δt is the time interval.

No target birth or death is considered, and the simulation does not model the sensor slew rate or settling time, local weather conditions, or lighting conditions of the objects being observed. The filter does model a 0.95 probability of detection for objects in the field of view and assumes a Poisson clutter model with a mean rate of 20 returns each epoch uniformly distributed in the FOV. The filter uses merging and pruning thresholds $U = 9$ and $T = 10^{-5}$ and a maximum of 3000 GMM components to maintain the Gaussian mixture approximation.

3.2.2. Results and Discussion

As in the linear model, the first case considers the use of the same TIF values for both TOIs and non-TOIs, setting $\tau_j = 1$. The 3D position error is computed for each object, by taking the maximum weighted component mean for each as the estimate and computing the Euclidean distance to the true value. Individual object errors are plotted in Figure 5(a) as gray lines, along with the average over all 1118 objects in black, and averages for TOIs and non-TOIs in blue and red respectively. In the case where objects are all given the same priority, the average position errors for TOIs and non-TOIs do not show much distinction, as expected. On the right, Figure 5(b) provides the results for the case where the non-TOI priority has been reduced significantly, to a value of 0.01. This choice is made instead of setting $\tau_{\text{non}} = 0$ to ensure non-TOIs still yield some information gain and produce valid sensor tasks in the case that no TOIs are visible to a given sensor at some point in time. The figure does show a sharp difference between the average position errors for TOIs and non-TOIs, discussed in further detail below.

Figure 6 provides a breakdown of the average position errors for TOIs and non-TOIs for all cases considered. On the left, the figure shows the best TOI errors are achieved for the case $\tau_{\text{non}} = 0.01$, and the final average position errors have converged to several hundred meters, approximately the level of the prescribed measurement noise. The $\tau_{\text{non}} = 0.1$ case similarly achieves excellent average position errors for TOIs by the final time. Both the $\tau_{\text{non}} = 0.5$ and $\tau_{\text{non}} = 1$ cases perform reasonably well, but have not converged at the final time as errors are still decreasing. On the right, the average non-TOI errors demonstrate that for the low TIF value, the system ignores them significantly as compared to the equal priorities case. The equal priorities case produces roughly the same position errors for both TOIs and non-TOIs as expected. Another noteworthy point from this figure is that the $\tau_{\text{non}} = 0.1$ case tracks both the best case TOI and non-TOI errors pretty well, offering an effective way to improve TOI tracking without sacrificing other objects all together.

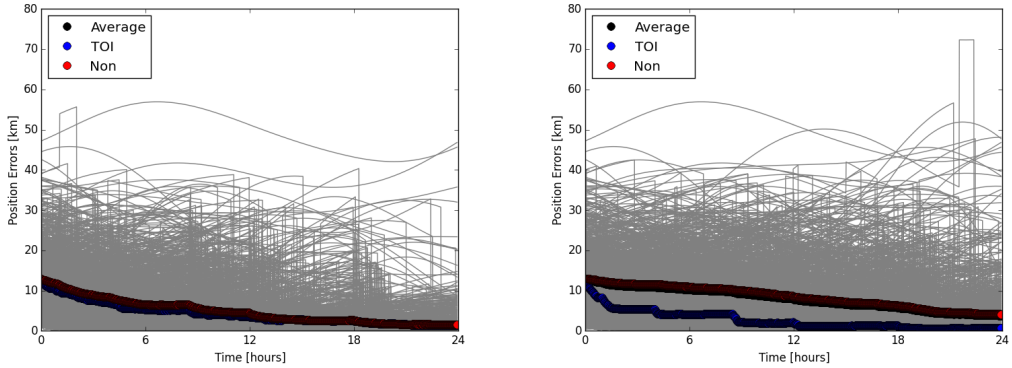

(a) Position Errors $\tau_{\text{non}} = 1$ (b) Position Errors $\tau_{\text{non}} = 0.01$

Fig. 5. Orbital Debris Position Errors

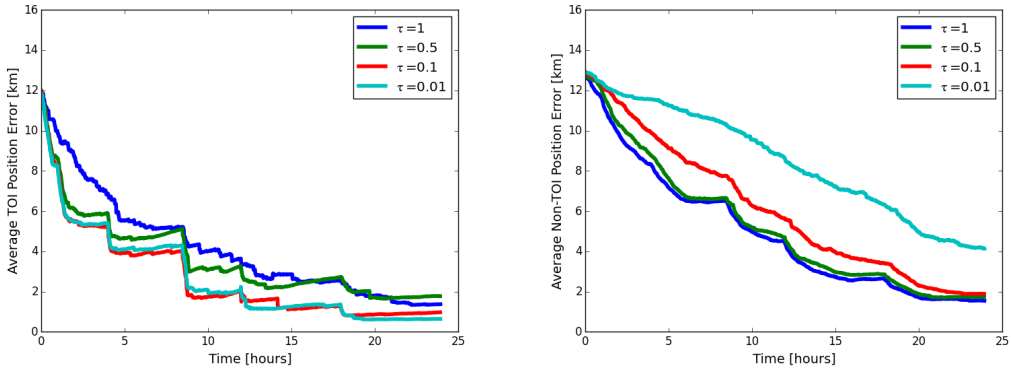

(a) TOI Position Errors (b) Non-TOI Position Errors

Fig. 6. Orbital Debris TOI and Non-TOI Errors

Table 4. Results Summary

Non-TOI Priority	Objects Detected	TOIs Detected	Number of Meas (TOI)	Percent Sensor Time TOIs
$\tau = 1$	1103/1118	111/112	6261 (616)	17.2
$\tau = 0.5$	1098/1118	109/112	6553 (852)	23.9
$\tau = 0.1$	1094/1118	111/112	7053 (1684)	46.8
$\tau = 0.01$	1045/1118	112/112	7166 (2946)	80.6

Table 4 provides a final summary of the results for each case, listing the total number of objects detected, number of TOIs detected, total number of measurements, and those for TOIs, and the percentage of sensor time spent tracking TOIs. The last two columns are different in that the number of measurements includes the effects of missed detections, and also cases where multiple objects are detected in a single scan. By contrast, the sensor time values indicate the percentage of sensor scans that are scheduled to observe at least one TOI, regardless of whether it is detected, or whether other objects are present in the field of view. As clear in the table, decreasing non-TOI priorities produces a corresponding increase in the amount of sensor time tracking TOIs. In addition, the $\tau_{\text{non}} = 0.01$ case is the only one to observe all TOIs at least once during the course of the simulation. The equal probabilities case produces the most detected objects overall. Interestingly, decreasing non-TOI priorities produces higher numbers of measurements, even

for non-TOIs. The likely explanation is that these cases favor observing multiple targets together, since single low priority objects would be less likely than multiple low priority objects to yield an expected information gain higher than scans with at least one TOI present.

4. CONCLUSIONS AND FUTURE WORK

This paper presented the theory and application of an information theoretic tasking scheme incorporating target priorities in the reward function. Through simulated test cases, the method was demonstrated to produce additional measurements and improved state estimates for higher priority targets. Many future applications are possible for this work, from implementing different sensor tasking modes designed to focus on certain objects, to automating tracking necessary to refine estimates of conjuncting objects and their associated probability of collision.

5. ACKNOWLEDGMENTS

The authors would like to acknowledge the support of the Cooperative Research Centre for Space Environment Management (SERC Limited) through the Australian Government's Cooperative Research Centre Programme.

6. REFERENCES

1. B. Ristic, B.-N. Vo, and D. Clark. A note on the reward function for phd filters with sensor control. *IEEE Transactions on Aerospace and Electronic Systems*, 47(2):1521–1529, 2011.

2. H. G. Hoang and B.-T. Vo. Sensor management for multi-target tracking via multi-bernoulli filtering. *Automatica*, 50(4):1135–1142, 2014.

3. Ronald P. S. Mahler. *Advances in Statistical Multisource-Multitarget Information Fusion*. Artech House, Norwood, MA, 2014.

4. Ronald P. S. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, Norwood, MA, 2007.

5. Ba-Tuong Vo, Ba-Ngu Vo, and Antonio Cantoni. Analytic implementations of the cardinalized probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 55(7):3553–3567, 2007.

6. M. Ulmke, Ozgur Erdinc, and Peter Willett. Gaussian mixture cardianlized phd filter for ground moving targets. In *Information Fusion (FUSION), 2007 10th International Conference on*, Quebec, Canada, July 2007.

7. Ba-Ngu Vo and Wing-Kin Ma. The gaussian mixture probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, 2006.

8. Simon J. Julier and Jeffrey K. Uhlmann. New extension of the kalman filter to nonlinear systems. In *Proc. SPIE 3068, Signal Processing, Sensor Fusion, and Target Recognition VI*, pages 182–193, Orlando, 1997.

9. Eric A. Wan and Rudolph Van Der Merwe. The unscented kalman filter for nonlinear estimation. In *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*, pages 153–158, Lake Louise, Alberta, October 2000.

10. Chee Sing Lee, Daniel E. Clark, and Joaquim Salvi. Slam with dynamic targets via single-cluster phd filtering. *IEEE Journal of Selected Topics in Signal Processing*, 7(3):543–552, June 2013.

11. Steven Gehly. *Estimation of Geosynchronous Space Objects Using Finite Set Statistics Filtering Methods*. PhD thesis, University of Colorado, Boulder, Boulder, CO, 2016.

12. Kyle J. DeMars, Robert H. Bishop, and Moriba K. Jah. Entropy-based approach for uncertainty propagation of nonlinear dynamical systems. *Journal of Guidance, Control, and Dynamics*, 36(4):1047–1057, July-August 2013.

13. Alfred O. Hero, Chris Kreucher, and Doron Blatt. Information theoretic approaches to sensor management. In A. O. Hero, D. A. Castanón, D. Cochran, and K. Kastella, editors, *Foundations and Applications of Sensor Management*, pages 33–57. Springer, New York, 2008.

14. Michael Beard, Ba-Tuong Vo, Ba-Ngu Vo, and Sanjeev Arulampalam. Sensor control for multi-target tracking using cauchy-schwarz divergence. In *Information Fusion (FUSION), 2015 18th International Conference on*, pages 937–944, Washington, D.C., July 2015.

15. Ronald Mahler. Multitarget sensor management of dispersed mobile sensors. In D. Grundel, R. Murphey, and P. Pardalos, editors, *Theory and Algorithms for Cooperative Systems*, pages 239–310. World Scientific Publishing Co., Singapoore, 2004.

16. Ronald P. S. Mahler. Sensor management with non-ideal sensor dynamics. In *Information Fusion (FUSION), 2004 7th International Conference on*, Stockholm, Sweden, June-July 2004.

17. Ronald P. S. Mahler. Unified sensor management using cphd filters. In *Information Fusion (FUSION), 2007 10th International Conference on*, August 2007.

18. Byron Tapley, Bob Schutz, and George Born. *Statistical Orbit Determination*. Elsevier Academic Press, New York, first edition, 2004.

19. David A. Vallado. *Fundamentals of Astrodynamics and Applications*. Microcosm Press, Hawthorne, CA, third edition, 2007.

20. Walter J. Faccenda, David Ferris, C. Max Williams, and Dave Brisnehan. Deep stare technical advancements and status. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, HI, September 2003.

21. William M. Folkner, James G. Williams, Dale H. Boggs, Ryan S. Park, and Petr Kuchynka. The planetary and lunar ephemerides de430 de431. IPN Progress Report 42-196, Jet Propulsion Laboratory, California Institute of Technology, http://ipnpr.jpl.nasa.gov/progress_report/42-196/196C.pdf, February 2009.

22. Nikolaos K. Pavlis, Simon A. Holmes, Steve C. Kenyon, and John K. Factor. An Earth gravitational model to degree 2160: EGM2008. In *Proceedings of the European Geosciences Union General Assembly*, Vienna, Austria, April 13-18 2008.

23. Chia-Chun Chao. *Applied Orbit Perturbation and Maintenance*. The Aerospace Press, El Segundo, CA, 2005.