AN EXPERIMENTAL IMPLEMENTATION OF A PARTICLE-BASED DYNAMIC SENSOR STEERING METHOD FOR TRACKING AND SEARCHING FOR SPACE OBJECTS

Tyler A. Hobson, I. Vaughan L. Clarkson

School of Information Technology and Electircal Engineering
The University of Queensland, Qld 4072, AUSTRALIA

ABSTRACT

We present practical, experimental results for a system, driven by a particle filter, that dynamically steers a space surveillance sensor to track and search for resident space objects. In contrast to traditional Kalman-filter-based trackers, this system can exploit scheduled observations where the target is not found within the field of view. Furthermore, real-time observation-evaluation enables the system to immediately respond to these events by conducting a limited search. We describe the system and report the results of a recent field trial using a computer-controlled Raven-class electro-optical sensor to track objects using two-line element sets (TLEs) of various ages. Even for quite old TLEs - in some cases over six months old - the system demonstrates successful, automatic reacquisition.

Index Terms— Particle filters, Space Situational Awareness, Intelligent sensors

1. INTRODUCTION

More than five decades of space-faring has resulted in a large population of man-made objects in Earth orbit [1]. Due to concerns that these objects may undergo destructive collisions with important space assets, several agencies endeavour to track as many objects as possible with the goal of maintaining Space Situational Awareness [2]. Nonetheless, some of these objects are too difficult to routinely observe or are perturbed too unpredictably to be tracked effectively and may on occasion become lost [3, 4].

Unless it is observed frequently, an object's state estimation error grows with time because of perturbations from imperfect modelling of the orbital dynamics and for active spacecraft, manoeuvring and station-keeping [5]. Gaussian recursive filtering techniques are commonly used [5–8] to obtain a value for this error bound by estimating the covariance of the state error p.d.f. This value is used to determine how often an object is to be observed to ensure the target has a high likelihood of residing within a sensor's limited Field of View (FOV). Nonetheless, a high probability of detection cannot always be guaranteed and Gaussian techniques do not have the capacity to store and utilise the information gained when the object is scheduled for observation but is not found in the FOV.

To capture this information, about where the target is not, we look to the Particle Filter (PF) [9]. The PF's point mass p.d.f. representation is capable of capturing a regional reduction in probability resulting from a failure to observe the target. Furthermore, the PF is capable of representing the shape of the p.d.f. to an arbitrary level of fidelity which provides valuable insight as to where to look next to find the target.

This work was supported by AFOSR grant FA9550-10-1-0493.

PFs have become popular over the last decade as a result of recent algorithmic innovations and modern computational power [9, 10]. Nonetheless, the majority of contemporary applications for tracking focus on sequential update assuming successful observations continue to arrive with certainty [9, 11], or else observations are discarded completely. Whilst there has been sporadic consideration of how to update a PF utilising the information gained when a target is not in the FOV [12–14], to the authors' knowledge, this principle has not been applied to the classical statistical tracking problem. Recently, the authors proposed a method [15] for modifying a PF to exploit information regarding the presence and the absence of a target within the FOV to provide a sensor with the intelligence to track and, if necessary, search for its target. SSA was used as an example application.

We present the results of an experimental implementation of the PF-based track-and-search method for controlling an electro-optical sensor with finite FOV. Several GPS satellites were observed over a number of nights using a set of initialisation data ranging in accuracy. The sensor, a Raven-class telescope, was dynamically steered toward a region of high probability using the PF's estimated p.d.f. which is updated by each observation. If the target is not observed and tracked, the system uses its estimated p.d.f. to direct the sensor toward the next most likely region in real time. The experimental results show that the method is not only effective, but that it is also practical to implement using existing technology.

In Section 2, we provide an overview of the PF-based track-and-search method proposed in [15]. The experimental implementation of the method is described in Section 3. Section 4 presents the results of this initial experiment and Section 5 offers some concluding remarks. Section 6 acknowledges the many people without whose support, the field trial would not have been possible.

2. PF-BASED SENSOR STEERING

Particle Filtering [9, 16, 17] approximates the underlying state p.d.f. to an arbitrary level of accuracy, using a point mass distribution. The points are referred to as particles and their corresponding masses as weights. Each particle $x^i; i = 1, 2, \ldots, N$ is drawn and weighted by w^i, where $\sum_{i=1}^{N} w^i = 1$, to characterise the posterior density $p(\mathbf{x}_{0:k} \mid \mathbf{z}_{0:k})$. The discrete-time sequence of state particle sets and observation vectors from time 0 to k are denoted by $\mathbf{x}_{0:k}$ and $\mathbf{z}_{0:k}$ respectively. Hence the approximation of the posterior distribution is made by

$$p(\mathbf{x}_{0:k} \mid \mathbf{z}_{0:k}) \approx \sum_{i=1}^{N} w_k^i \delta(\mathbf{x}_{0:k} - \mathbf{x}_{0:k}^i) \quad (1)$$

where $\delta(\cdot)$ is the Dirac delta function.

In common with other recursive filters, PFs execute both a predictive time update and a corrective measurement update at each time step or epoch k. The predictive step results in $\hat{p}(\mathbf{x}_{0:k} \mid \mathbf{z}_{0:k-1})$ and the corrective step in $\hat{p}(\mathbf{x}_{0:k} \mid \mathbf{z}_{0:k})$. Thereafter resampling may be necessary to encourage particles to reside in high-likelihood regions and discourage a situation where all but a few particles have non-trivial weights [9].

The question of how to aim the sensor is not often considered, it being perhaps tacitly assumed that a good enough track should be maintained to ensure that the target is always observed when the sensor is aimed towards it. The principal adaptation here is to give explicit consideration to how the sensor is steered and to how the particles and weights should be updated if the object is not within the FOV. As proposed in [15] in greater detail, this adaptation is achieved by incorporating a search component between the prediction and measurement-update steps of the classical PF. The search component comprises of two elements which consist of a sensor steering routine and a measurement evaluation method.

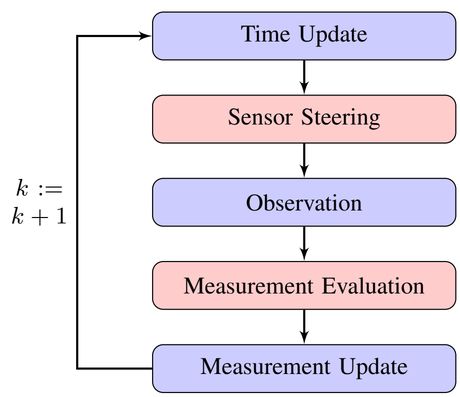

Fig. 1. PF sequential update process with search elements

The sensor steering routine utilises an objective function $J(\mathbf{\Psi})$, to determine the most appropriate sensor pointing angles. The steering routine is incorporated immediately proceeding the PF time update step to ensure $\hat{p}(\mathbf{x}_{0:k} \mid \mathbf{z}_{0:k-1})$ is computed prior to the evaluation of steering parameters $\mathbf{\Psi}$, such as a set of proposal pointing angles. Utilising $\hat{p}(\mathbf{x}_{0:k} \mid \mathbf{z}_{0:k-1})$, the objective function may evaluate $\mathbf{\Psi}$ with metrics such as p.d.f. variance reduction or the number and weight of particles in the FOV. Hence, in Fig. 1, the steps shaded in blue represent the actions of a traditional Kalman-filter-based tracker, to which we explicitly add consideration of the steps shaded in red: sensor steering and measurement evaluation.

The measurement evaluation method occurs prior to the PF update step to inform the filter how it should treat the latest observation. The measurement evaluation is based on performance parameters, $\mathbf{\Omega}$ which may include information such as association confidence and inclement weather detection. These parameters are collected during the observation enabling the system to evaluate if the target was in the FOV, compute the probability of Type I & II errors occurring during the measurement and thereafter select an appropriate proposal density $q(\cdot)$ to apply during the measurement update. If the target is observed in the FOV, $q(\cdot)$ should bolster the particle weights surrounding the target. If the target is not observed, the weights will be reduced to discourage further observation of this region of the p.d.f. In either case, the estimated probability of Type I & II errors may be used to appropriately weight the particles in the FOV. In circumstances where certain other error conditions arise, such as when cloud obstructs visibility, observations should be dismissed entirely.

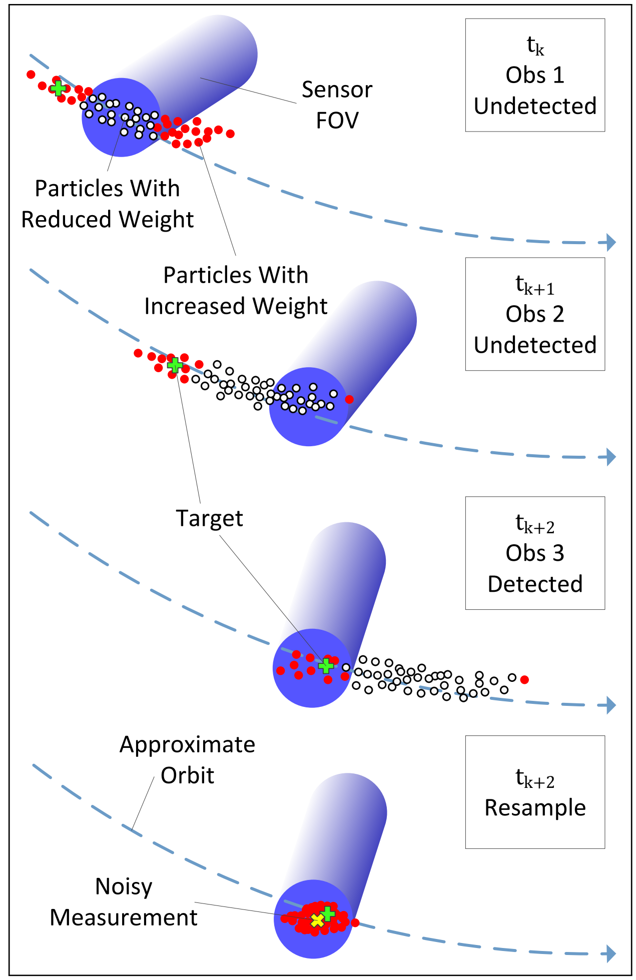

Fig. 2. Depiction of the PF-based search strategy

The iteration of this process, as shown diagrammatically in Fig 2, enables the filter to judiciously select a new pointing angle based on the previous observation, irrespective of whether the target was located in the FOV. As a consequence the filter is capable of dynamically steering the sensor to automatically track and search as required. If the p.d.f. is well conditioned such that the sensor has a high likelihood of seeing the target, the system will track the target. If by chance the target is not found in the FOV, the system has the capability to execute a directed search until it achieves reacquisition.

3. EXPERIMENTAL METHOD

3.1. System Architecture

Implementation of the sensor control strategy, as described in [15] and reviewed in the previous section, required the integration of a number of disparate systems. The primary control program named Space PARticle Search Evaluation (SPARSE) was developed in MATLAB. SPARSE monitors and controls all aspects of the search-track strategy. In addition to implementing the modified PF, SPARSE employs physical models including sensor, planetary, lunar, solar illumination and space object orbit propagation models for

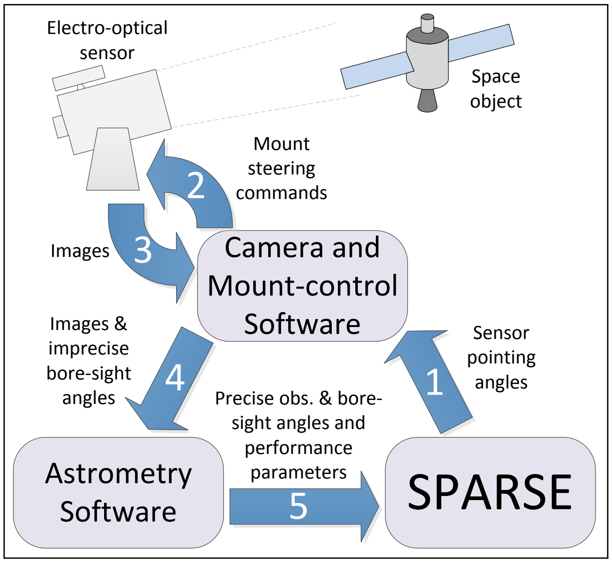

Fig. 3. Architecture of experimental system

target visibility prediction and to obtain the information necessary to run the modified PF.

SPARSE is designed, but not limited, to interface with Raven class electro-optical equipment, to which controlled access was graciously offered by the US Air Force Research Laboratory (AFRL) to implement the experiment as part of an ongoing collaboration. The Raven class system comprises of a telescope, a computer-controlled mount, astrometry software capable of measuring a target's observation angles with a standard deviation of one arcsecond as well as camera and mount-control software. The high-level test architecture is shown in Fig. 3. The numbering indicates the order in which information is passed throughout the system in a dynamic loop.

3.2. System Configuration

At the beginning of a loop, SPARSE applies its physical models to its latest object state estimates to determine if any objects are currently visible to the sensor. Thereafter, the list of visible objects is used by SPARSE to choose its current target according to programmable criteria. For this experiment, a target was chosen at random. After propagating the target's particles to the current observation epoch, SPARSE employs the sensor steering routine to produce the steering commands to be sent to the camera and mount control software. Whilst the authors have assessed a range of objective functions to steer the sensor, the most effective function was chosen for presentation in this paper. The objective function chosen to produce the results presented in Section 4, evaluates a discrete set of pointing angles constrained to align the sensor's bore-sight with each of the PF's particles. The i^{th} particle is selected by evaluating which of the N pointing angles maximises the likelihood of detection, as estimated by the particle distribution. The resulting objective function is described by,

$$J\left(\boldsymbol{\Psi}^{i}\right)=\sum_{s\in\mathcal{S}^{i}} w_{k-1}^{s}, \tag{2}$$

where $\mathcal{S}^i$ is the set of indices for particles that will fall within the sensor's FOV when the sensor is aimed at the i^{th} particle.

Once the telescope is steered and the images returned, the astrometry software performs astrometric correlation on the images to obtain precise pointing angles. In addition, it evaluates the observation performance metrics $\boldsymbol{\Omega}_k$ and returns the results to SPARSE. The primary performance metrics returned to SPARSE includes a level of data association confidence concerning any objects within the FOV and a subjective assessment of any occlusion due to inclement weather, according to the number of stars visible in each image.

Once SPARSE is supplied with $\boldsymbol{\Omega}_k$, it awards each observation with one of the following three states:

1. WITHIN FOV - the target was observed in the FOV,
2. OUTSIDE FOV - the observation is valid but the target was not observed in the FOV,
3. INVALID - the probability of Type I & II errors is too high.

If an observation is classified as WITHIN FOV, a proposal density resulting in a classical PF Update step is calculated [15]. The particle weights are scaled according to

$$w_k^i \propto w_{k-1}^i p\left(\mathbf{z}_k \mid \mathbf{x}_k^i\right), \tag{3}$$

which results in an increase in relative particle weights and/or particles around the observation, scaled by the assumed observation noise. An OUTSIDE FOV results in a proposal density that nullifies all weights of the particles within the FOV, such that

$$w_k^s = 0; \forall s \in \mathcal{S}^i. \tag{4}$$

The INVALID state is awarded in the event reliable information cannot be obtained regarding the presence or absence of the target within the FOV. As a result, the filter is not updated and the observation is discarded as though it was never performed.

3.3. Test Procedure

With the cooperation of AFRL at Maui, under the leadership of Dr. Kim Luu and with the support of Pacific Defence Solutions (PDS), a field trial was conducted on 22–27 October, 2013. The trial made use of a Raven-class telescope at Kihei. To test the ability of the proposed method to reacquire objects that are temporarily lost, Two-Line Elements (TLEs) of various ages were used as prior information. Older TLEs yield less accurate information about the current position of an object and so are expected to trigger search behaviour more often. The aim is to see whether a track can be reacquired even with quite old TLEs. Since regularly updated orbital element sets and truth data were desirable for the test, GPS satellites were used as the targets. The objects that were used to produce the results presented in Section 4 are listed in Table 1 using their USSTRATCOM catalogue identification number. TLEs for each

Table 1. GPS satellite target list for experiment

US Cat. ID	Name
26360	NAVSTAR 47 (USA 150)
32711	NAVSTAR 62 (USA 201)
25030	NAVSTAR 44 (USA 135)
35752	NAVSTAR 64 (USA 206)
22014	NAVSTAR 26 (USA 83)
32260	NAVSTAR 60 (USA 196)

object were obtained from [18] that were approximately 0, 50, 100, 150 and 200 days old.

Once the TLEs are loaded into SPARSE, SPARSE uses its internal SPG4 propagator [19] to independently decode each TLE and turn them into state particles with equal weight. The particles are then propagated forward in time to the first observation epoch.

In accordance with the method proposed in [15], each object is propagated with sufficient process error to accommodate the dynamics left unmodelled by the SGP4 propagator. For this experiment, whilst it was anticipated that each object should experience similar levels of perturbation, it was desirable to learn how the system would behave when the process error of some of the objects is inflated. For this reason, the particles for every second object, with respect to Table 1, were supplied with at least 10 times more process error than was necessary for robust tracking.

Each set of particles is used independently to aim the electro-optical sensor. Each object and each TLE was scheduled one set of five contiguous observations each night. The result of each observation affected in real-time how the sensor was pointed during the next observation.

Table 2. Observation States Recorded During Field Test

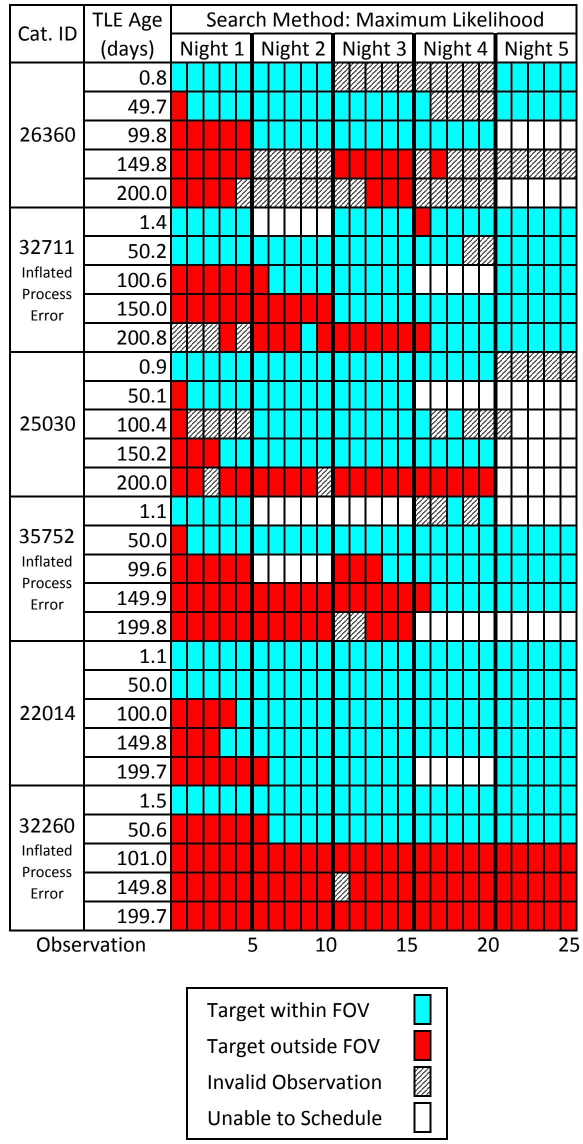

4. RESULTS

Table 2 displays the results of running the experimental system over a period of 5 nights. Operational constraints limited sensor availability such that there was not always enough time to observe every historical TLE each evening. This is indicated by the blank squares in Table 2. In spite of these constraints and periods of cloud cover causing invalid observations, many examples of target reacquisition were recorded. Scanning left to right along the rows of Table 2, it is seen that there are 15 instances in which the target was initially outside the FOV (red) but was eventually reacquired (blue) through the automatic search conducted by SPARSE.

Table 2 shows a clear correlation between the age of the TLE and the number of observations necessary to observe or reacquire a target. There is also a noticeable increase in the number of days necessary to reacquire targets with inflated process error. The extreme case is object 32260 whose TLEs of 100 days or older contained too much error to allow reacquisition in the time allowed. Nevertheless, in most cases, the systematic search is able to reacquire the target and then consistently track the object through the remaining scheduled observations.

5. CONCLUSION

The results of the field trial demonstrate the ability of a modified particle filter to dynamically control a sensor to track and search for targets under surveillance. Unlike conventional tracking systems that would in most of the test cases classify the object as lost, SPARSE used subsequent observations to search and in many cases reacquire the target in a judicious manner.

Whilst not all targets were reacquired for all TLEs, the authors believe that, had the trial lasted a little longer and encountered fewer interruptions, the few remaining objects may have been reacquired and tracked.

In conclusion, we have demonstrated a practical system for the dynamic steering of space sensors which not only tracks RSOs but automatically reacquires objects that are temporarily lost. Remarkably, these preliminary results suggest it is feasible both theoretically and practically to routinely reacquire objects using TLEs that are more than 3 months old.

6. ACKNOWLEDGEMENTS

The authors would like to thank Kim Luu and Chris Sabol of AFRL for their guidance, sharing their experience and for their efforts ensuring the field trials became a reality.

The authors would also like to acknowledge the team at PDS, in alphabetical order, Randy Cortez, Keric Hill, Laurie Hozaki, Kawai Kuluhiwa, Daron Nishimoto, Vicky Sensano and Paul Sydney for hosting the experiment, offering their technical expertise and supporting the operation of the experimental system and its many components.

7. REFERENCES

[1] Nickolay N. Smirnov, Ed., *Space Debris: Hazard Evaluation and Mitigation*, Earth Space Institute. CRC Press, London, New York, 2002.

[2] David A. Vallado and Jacob D. Griesbach, "Simulating space surveillance networks," in *Proc. AAS/AIAA Astrodynamics Specialists Conf.*, 2012.

[3] Tom Kelecy and Moriba Jah, "Analysis of orbit prediction sensitivity to thermal emissions acceleration modeling for high area-to-mass ratio (HAMR) objects," in *Proc. the Advanced Maui Optical and Space Surveillance Technologies Conf., (AMOS)*, 2009.

[4] Jayant Sharma, Grant H. Stokes, Curt von Braun, George Zollinger, and Andrew J. Wiseman, "Toward operational space-based space surveillance," *Lincoln Laboratory Journal*, vol. 13, no. 2, pp. 309–334, 2002.

[5] Byron D. Tapley, Bob E. Schutz, and George H. Born, *Statistical Orbit Determination*, Elsevier Academic Press, Amsterdam; Boston, 2004.

[6] Keric Hill, Paul Sydney, Kris Hamada, Randy Cortez, Kim Luu, Moriba Jah, Paul W. Schumacher Jr., Michael Coulman, Jeff Houchard, and Dale Naho'olewa, "Covariance-based network tasking of optical sensors," in *Proc. AAS/AIAA Space Flight Mechanics Meeting*, 2010.

[7] John Seago, Jacob Griesbach, James Woodburn, and David A. Vallado, "Sequential orbit-estimation with sparse tracking," in *Proc. AAS/AIAA Space Flight Mechanics Conf.*, 2011.

[8] David A. Vallado and Wayne D. McClain, *Fundamentals of Astrodynamics and Applications*, Space Technology Library. Springer, New York, 3rd edition, 2007.

[9] Branko Ristic, Sanjeev Arulampalam, and Neil Gordon, *Beyond the Kalman Filter*, Artech House, 2004.

[10] J. H. Kotecha, Jianqui Zhang, Yufei Huang, T. Ghirmai, M. F. Bugallo, and J. Miguez, "Particle Filtering," *IEEE Signal Processing Magazine*, vol. 20, no. 5, pp. 19–38, 2003.

[11] O. Cappe, S. J Godsill, and E. Moulines, "An overview of existing methods and recent advances in sequential Monte Carlo," *Proceedings of the IEEE*, vol. 95, no. 5, pp. 899–924, 2007.

[12] Henry R. Richardson, Lawrence D. Stone, W. Reynolds Monach, and Joseph H. Discenza, "Early maritime applications of particle filtering," in *Proc. SPIE 5204, Signal and Data Processing of Small Targets*, 2003.

[13] Frederic Bourgault, Tomonari Furukawa, and Hugh F. Durrant-Whyte, "Coordinated Decentralized Search for a Lost Target in a Bayesian World," in *Proc. 2003 IEEE/RSJ International Conf. on Intelligent Robots and Systems*, 27-31 Oct 2003, vol. 1, pp. 48 – 53.

[14] Tomonari Furukawa, Frederic Bourgault, Benjamin Lavis, and Hugh F. Durrant-Whyte, "Recursive Bayesian Search-and-Tracking Using Coordinated UAVs for Lost Targets," in *Proc. IEEE International Conf. on Robotics and Automation*, 15-19 May 2006, pp. 2521 – 2526.

[15] Tyler A. Hobson and I. Vaughan L. Clarkson, "A Particle-based Search Strategy for Improved Space Situational Awareness," accepted in Asilomar Conf. on Signals Systems, and Computers, 3-6 Nov 2013.

[16] Volkan Cevher and James H. McClellan, "Fast initialization of particle filters using a modified Metropolis-Hastings algorithm: mode-hungry approach," in *Proc. IEEE Acoustics, Speech, and Signal Processing Conf., (ICASSP)*, 2004, vol. 2, pp. 129–32.

[17] Walter R. Gilks and Carlo Berzuini, "Following a moving target — Monte Carlo inference for dynamic Bayesian models," *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, vol. 63, no. 1, pp. 127–146, 2001.

[18] USSTRATCOM, "Space Track," 2004, [Online]. Available: www.space-track.org.

[19] David A. Vallado, Paul Crawford, Richard Hujsak, and T. S. Kelso, "Revisiting Spacetrack Report NO.3," in *Proc. AIAA/AAS Astrodynamics Specialist Conf.*, Aug 2006, number AIAA 2006-6573.