

Autonomous space object catalogue construction and upkeep using sensor control theory

Nicholas Moretti

Inovor Technologies

Mark Rutten and Travis Bessell

National Security and Intelligence, Surveillance and Reconnaissance Division

Defence Science and Technology Group

Brittany Morreale

United States Air Force

ABSTRACT

The capability to track objects in space is critical to safeguard domestic and international space assets. Infrequent measurement opportunities, complex dynamics and partial observability of orbital state makes the tracking of resident space objects nontrivial. It is not uncommon for human operators to intervene with space tracking systems, particularly in scheduling sensors.

This paper details the development of a system that maintains a catalogue of geostationary objects through dynamically tasking sensors in real time by managing the uncertainty of object states.

As the number of objects in space grows the potential for collision grows exponentially. Being able to provide accurate assessment to operators regarding costly collision avoidance manoeuvres is paramount; the accuracy of which is highly dependent on how object states are estimated. The system represents object state and uncertainty using particles and utilises a particle filter for state estimation. Particle filters capture the model and measurement uncertainty accurately, allowing for a more comprehensive representation of the state's probability density function.

Additionally, the number of objects in space is growing disproportionately to the number of sensors used to track them. Maintaining precise positions for all objects places large loads on sensors, limiting the time available to search for new objects or track high priority objects. Rather than precisely track all objects our system manages the uncertainty in orbital state for each object independently. The uncertainty is allowed to grow and sensor data is only requested when the uncertainty must be reduced. For example when object uncertainties overlap leading to data association issues or if the uncertainty grows to beyond a field of view. These control laws are formulated into a cost function, which is optimised in real time to task sensors.

By controlling an optical telescope the system has been able to construct and maintain a catalogue of approximately 100 geostationary objects.

1 INTRODUCTION

Space Situational Awareness (SSA) is the ability to detect, track and predict the location of resident space objects (RSOs). Modern society, industry and military rely on space assets to perform a number of critical tasks, the failure of one of these space assets could be catastrophic. Collecting precise SSA data and constructing accurate products from them is paramount to safe guarding space assets from collisions or from malicious activity. With the emergence of "Space 2.0" and space becoming more accessible, the volume of data to be collected is ever increasing, further complicating the SSA problem.

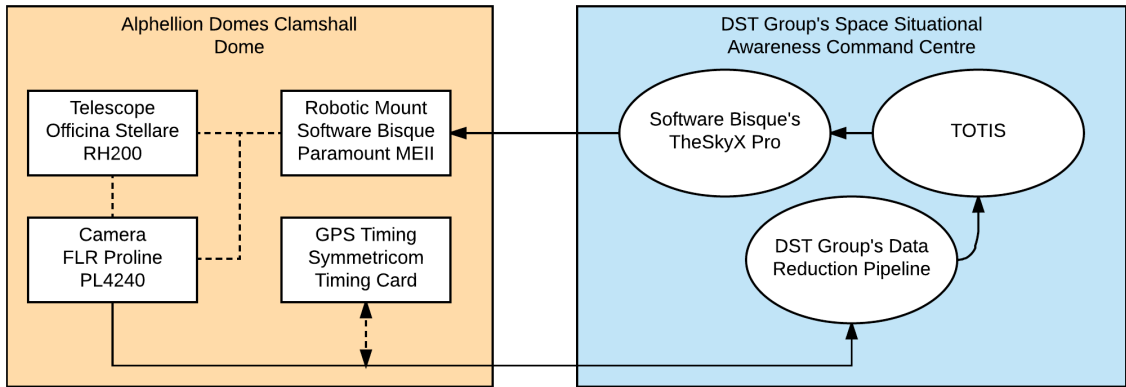


Fig. 1. DST Group's SSA facilities used with TOTIS

Typical SSA systems do not have direct control of their associated sensors. Generally, at the start of an observation period (e.g. a night for an optical sensor), the sensor is tasked for the entire period. This results in the sensor being unresponsive to dynamic changes in the catalogue or space environment, requiring significant human intervention.

Tracker of Things in Space (TOTIS) was developed to be a completely autonomous system to construct, refine and maintain a catalogue of RSOs in real time. TOTIS was developed as a research tool for the Defence Science and Technology Group (DST Group), part of the Australian Department of Defence, with the aim of being able to construct, refine and maintain a catalogue of geosynchronous like RSOs. TOTIS's main objectives are to, in real time, maintain the uncertainty of RSOs within limits to avoid data association errors and/or cross tagging. The detections found by TOTIS are passed to a different tool to construct high fidelity orbital state estimates off-line. TOTIS's aim is not to construct highly precise orbital estimates, but rather be a real-time enabler for another tool to do this.

TOTIS employs state estimation techniques not generally used within SSA to more accurately represent object state and its uncertainty, allowing for greater accuracy outputs and assessments of RSOs to be produced. TOTIS uses dynamic sensor control to manage the uncertainty in objects individually so they can be re-acquired, while still surveying segments of the sky for new RSOs.

This paper details the development of TOTIS and an experimental trial validating its operation. Section 2 gives an overview of the system, Section 3 reviews the sensor control methodology, Section 4 and Section 5 detail the state estimation and associated initial orbit determination techniques respectively, and Section 6 describes a multiple day validation campaign.

2 SYSTEM OVERVIEW

DST Group is developing a research facility for Space Situational Awareness. This facility has been used in a number of previous SSA research activities [1][2], and was used to develop and demonstrate TOTIS. The facility includes a robotic electro-optical SSA sensor and consumer grade computing facilities; shown in Fig. 1. The electro-optical SSA sensor is an Officina Stellare RH200 telescope paired with an FLI Proline PL4240 camera. This sensor operates in the visible wavelength and has a field of view of approximately 2.6 degrees with 4.6 arcseconds per pixel. The sensor is mounted to a Software Bisque Paramount MEII robotic mount, inside an Aphelion Domes 7ft clamshell dome. A Symmetricom (now Microsemi) GPS timing card is connected to the camera shutter to precisely time-stamp the opening and closing of the camera shutter. The sensor hardware is located at Edinburgh, South Australia.

Located remotely from the sensor hardware is DST Group's Space Situational Awareness Control Centre. This facility contains a computer running TOTIS and associated software peripherals. The computer is consumer grade and includes a Nvidia GTX780 Graphics Processing Unit for performing computation in parallel. Software Bisque's TheSkyX Pro is used to interface, command and control the robotic mount and camera. TOTIS interfaces with the telescope via TheSkyX Pro using Software Bisque's scripting

function. Captured photos are passed through DST Group’s custom developed data reduction pipeline to extract satellites and perform astrometry. TOTIS then ingests the data to perform its SSA tasks.

TOTIS can be split into four main components, sensor control, state estimation, data association and initial orbit determination. Each component, excluding data association, is described in a section below. The data association method used within TOTIS is Joint Probabilistic Data Association. It’s application and performance to SSA has previously been researched by the authors [3], as such it is not detailed in this paper. An outline of the sensor control methodology is presented with more detail in further work by the authors in [4].

3 SENSOR CONTROL

TOTIS was developed to be a real-time system to maintain a catalogue of RSOs. To enable this it was determined that there were two main needs:

1. The system must be able to reacquire RSOs with ease, and
2. The system must be able to associate detections to the catalogue minimising data association errors and cross tagging.

These high-level needs were derived into two requirements for the sensor control methodology

1. an RSO’s uncertainty will be reduced by a measurement (if possible) before its uncertainty extends beyond a sensor’s FOV, and
2. an RSO’s uncertainty will be reduced by a measurement (if possible) before its uncertainty overlaps another RSO’s uncertainty in measurement space.

In contrast to typical SSA systems, TOTIS has real-time control over its sensors. This allows for the above requirements to be formulated into cost functions that are optimised in real-time directly before each time step. TOTIS can immediately react to changes in environment and catalogue, such as weather resulting in a missed observation, or detection of a new RSO. If all of the object refinement requirements are met, then the surplus sensor resources are used to perform surveillance in order to search for new objects.

The control methodology used within TOTIS is highly flexible. As the cost functions are calculated for each RSO, the parameters can vary between RSOs, customising the sensor control methodology to each object. For example, the functions can be tailored to provide more frequent measurements on a subset of RSOs, or a maximum revisit time.

The detailed sensor control algorithms have been previously presented by the authors in [4].

4 STATE ESTIMATION

In TOTIS, the probability distribution representing an orbit is updated with new measurement information recursively using a recursive Bayesian state-space estimation method [5]. These methods are based around the definition of state-space equations. The process model, which defines how an object moves through space, and the measurement model, which defines the relationship between measurements and the underlying state. A general discrete-time process model can be written as

$$x_{k+1} = f(x_k, u_k), \tag{1}$$

where x_k is the state of the system at time k and u_k is a noise process that describes the uncertainty in the system. The process model defines the state transition probability density function (pdf) $p(x_{k+1}|x_k)$. The measurement model can be written as

$$z_k = h(x_k, v_k), \tag{2}$$

where z_k is the measurement at time k and v_k is a measurement noise process. The measurement model defines the measurement likelihood $p(z_k|x_k)$. A recursive Bayesian filter calculates the posterior pdf of

the state in terms of the state transition pdf and measurement likelihood

$$p(x_k|z_{1:k}) \propto p(z_k|x_k) \int p(x_k|x_{k-1})p(x_{k-1}|z_{1:k-1})x_{k-1}, \quad (3)$$

where $z_{1:k}$ represent the set of measurements from times 1 to k and it is assumed that the prior pdf, $p(x_0)$, is known. The Kalman filter is the ubiquitous example of a state-space estimation technique, which arises when the process and measurement models are linear and Gaussian. In that case (3) results in a recursion on the mean and covariance of the state.

In this work a particle filter is used to update the pdf of an orbit using measurement information [6]. Particle filters approximate the probability density function (pdf) of the state of an object as a set of weighted samples. Statistics derived from the distribution, such as the mean, can be calculated directly from the samples. If there are N particles with states given by $x^{(i)}$ and their normalised weights by $w^{(i)}$ ($\sum_{i=1}^N w^{(i)} = 1$), then the mean can be calculated as

$$\hat{x} = \sum_{i=1}^N w^{(i)} x^{(i)}. \quad (4)$$

A particle filter does not assume an analytic form for the state pdf (c.f. an unscented Kalman filter [5]) and there is no need to approximate the state-space equations (c.f. an extended Kalman filter), but this flexibility is a trade-off against a typically higher computational expense.

For a particle filter, the update equation (3) is encapsulated in a recursion on the particle weights. First a new sample, $x_k^{(i)}$ is generated from $x_{k-1}^{(i)}$ by sampling from a proposal distribution

$$x_k^{(i)} \sim q(x_k|x_{0:k-1}^{(i)}, z_{1:k}), \quad (5)$$

then the updated particle weight is given by

$$\tilde{w}_k^{(i)} \propto \frac{p(z_k|x_k^{(i)})p(x_k^{(i)}|x_{k-1}^{(i)})}{q(x_k^{(i)}|x_{0:k-1}^{(i)}, z_{1:k})} w_{k-1}^{(i)}. \quad (6)$$

The update is performed independently for each particle, then the weights are normalised

$$w_k^{(i)} = \frac{\tilde{w}_k^{(i)}}{\sum_j \tilde{w}_k^{(j)}}. \quad (7)$$

As measurement information is incorporated into the filter it is very common for the particle weights to become degenerate, that is only one particle has significant weight. This can be ameliorated by the choice of the proposal distribution (5) and through resampling. Resampling results in a set of particles with uniform weight and acts to replicate particles with high weight and remove particles with low weight [6]. The requirement for resampling is almost always unavoidable and is a key step in a particle filter recursion.

The particle filter equations simplify if the process equation is used as the proposal distribution, that is

$$x_k^{(i)} \sim p(x_k|x_{k-1}^{(i)}) \quad (8)$$

and resampling is performed at each iteration, which forces all of the weights to be equal. Then the weight equation (6) becomes

$$\tilde{w}_k^{(i)} \propto p(z_k|x_k^{(i)}), \quad (9)$$

resulting in the weight of each particle simply being the likelihood of the measurement for that particle.

For orbit propagation the prior distribution is very diffuse, while both the process noise and the measurement noise are typically very small. This is among the most challenging set of numerical requirements for a particle filter. Standard methods of defining efficient proposal distributions [7] are ineffective in this situation due to the small process noise relative to the extent of the prior. We have found that regularised particle filters perform well in this case [8]. Regularised particle filters use the particles to approximate a continuous density using kernel density estimation methods.

Particle filters require each particle to be passed through the process model and measurement model at each time step. This places significant demands on computational resources for a real-time system. For example, TOTIS represents each RSO with 50,000 particles, tracking 100 objects requires 5,000,000 orbital states to be propagated through orbital dynamics models each time step - a non-trivial task.

The ability to precisely represent the measurement and process functions within a particle filter allows for lower fidelity, less computationally demanding process models to be used by capturing the uncertainty in the model. Previous research [1] showed that the SGP4 algorithm can be used in combination with a GPU to propagate large numbers of orbital states in real-time. The SGP4 algorithm was trialed in TOTIS but proved problematic due to the transformation to and from mean orbital elements. The research in [1] showed that lower fidelity orbital dynamics models can be used within the particle filter framework as uncertainty can be more accurately captured. As such, two body orbital dynamics equations are used within TOTIS to propagate orbital state. As the aim of TOTIS is not to construct high-fidelity orbital states of RSOs we have found that the use of two body equations is acceptable. The use of simple dynamic equations ported to a GPU allows for TOTIS to operate in real-time while maintaining continuous tracks of geostationary objects.

5 INITIAL ORBIT DETERMINATION

Initial orbit determination (IOD) is the process by which a measurement uncorrelated to an RSO in the catalogue is introduced into the catalogue. The two-dimensional angles measurement information obtained from the optical sensor under-defines the six-dimensional orbital state. TOTIS captures multiple images in succession so angle rates can be (imprecisely) inferred, but range and range-rate are not available. Therefore, initial orbital state samples can not be directly obtained from the measurement distribution.

The method used to perform IOD in TOTIS is a Bayesian Markov-Chain Monte Carlo (MCMC) sampling technique [9], which is similar to the IOD method developed by Schneider [10]. In comparison to other constrained short-arc IOD methods, such as [11], MCMC allows us to incorporate the measurement data precisely and does not require an explicit approximation of angles rates, while providing a more flexible means of constraining the orbital solutions.

MCMC is a class of methods which generate samples from a target distribution which can be easily evaluated, but is difficult or impossible to sample from directly. This is achieved by forming a chain of samples from a *proposal* distribution defining the transitions in the chain that can be easily sampled from. A series of simple accept/reject rules on this chain result in a set of samples that are provably from the target distribution (in the limit).

Define the target probability distribution over the state space of x to be $\pi(x)$ and the MCMC proposal distribution to be $q(x|y)$, which is the probability of transition to state x from state y . The Metropolis-Hastings (MH) MCMC algorithm defines the accept/reject probability as

$$\alpha(x, y) = \min \left\{ 1, \frac{\pi(x)q(x|y)}{\pi(y)q(y|x)} \right\}. \quad (10)$$

Starting with an initial sample, x_0 , the MH sampler proceeds by sampling from $q(x|x_0)$ to give x' . This sample is then accepted or rejected so that the next sample in the chain, x_1 , is given by

$$x_1 = \begin{cases} x', & \text{if } \alpha(x_0, x') > \mu \quad (\text{accept move}) \\ x_0, & \text{otherwise} \quad (\text{reject move}), \end{cases} \quad (11)$$

where μ is a uniformly generated random number between 0 and 1. This chain of samples can be propagated indefinitely, resulting in a set of samples from the target distribution $\pi(x)$.

The proposal distribution that is used in this paper is a Gaussian random walk and this means that $q(x|y) = q(y|x)$ and so (10) simplifies to

$$\alpha(x, y) = \min \left\{ 1, \frac{\pi(x)}{\pi(y)} \right\}. \quad (12)$$

The MH MCMC algorithm in this case is just an evaluation of the target distribution at the current and proposed locations. If the probability at the proposed location is higher than at the current location, it

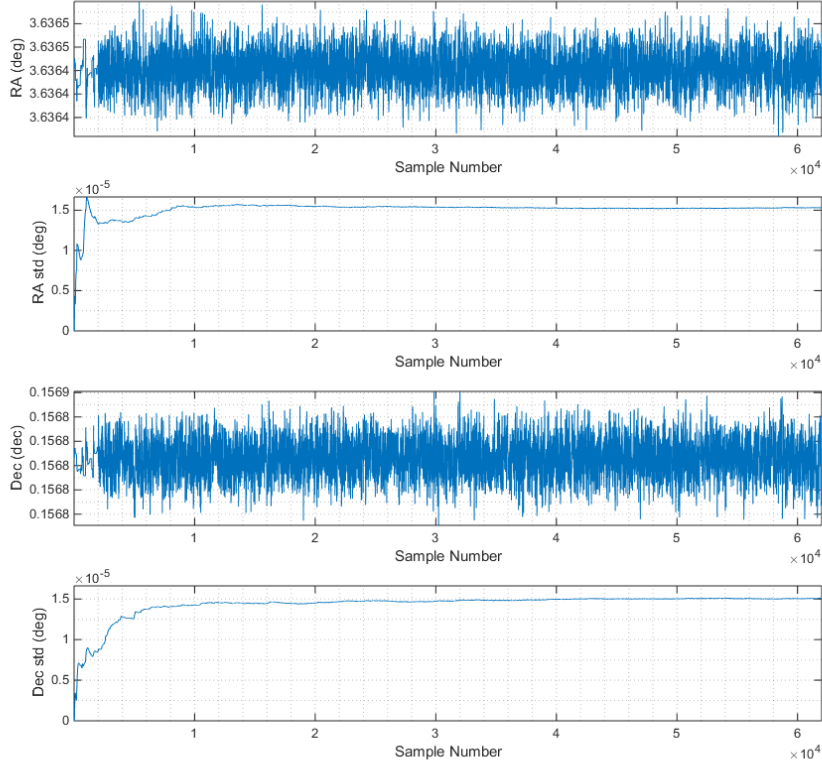


Fig. 2. Samples from the Markov Chain adapting to the target distribution in the Right Ascension and Declination of the object state-space. The first figure in the two pairs shows the samples from the chain, while the second shows the standard deviation of the proposal distribution.

is automatically accepted, otherwise it is randomly selected according to the ratio of probabilities, as in (11).

The initial orbit determination algorithm is supplied with a sequence of detections, which provide angles (right ascension and declination). Since the sensor has been characterised, there is knowledge about the accuracy of these measurements, supplied as a measurement noise covariance matrix, R . Given an orbital state, x , the likelihood of one of these measurements, z_i , can be written as the Gaussian probability

$$p(z_i|x) = \frac{1}{\sqrt{|2\pi R|}} \exp\left(-\frac{1}{2} [z_i - h_i(x)]^T R^{-1} [z_i - h_i(x)]\right), \quad (13)$$

where $h_i(x)$ is the measurement function relating the state x to the measurement z_i , as in (2). The IOD target distribution can then be written as

$$\pi(x|z_{1:N_m}) = p(x) \prod_{i=1}^{N_m} p(z_i|x), \quad (14)$$

where $\pi(x|z_{1:N_m})$ means the probability of x conditioned on the set of all N_m measurements to be used for IOD. The prior probability density, $p(x)$, encodes the constraints that we wish to impose on the orbits. In this case the prior probability is uniform over all eccentric orbits within predefined limits on the apogee and perigee of the orbits and zero outside of those limits. It is trivial to use alternative constraints, such as the eccentricity and semi-major axis limits [11].

The main challenge in implementing an efficient MH MCMC algorithm is the choice of the proposal distribution, $q(x|y)$. For example, if the proposal results in small jumps compared to the size of the target distribution, then the proportion of accepted jumps will be very high, but the chain will take a

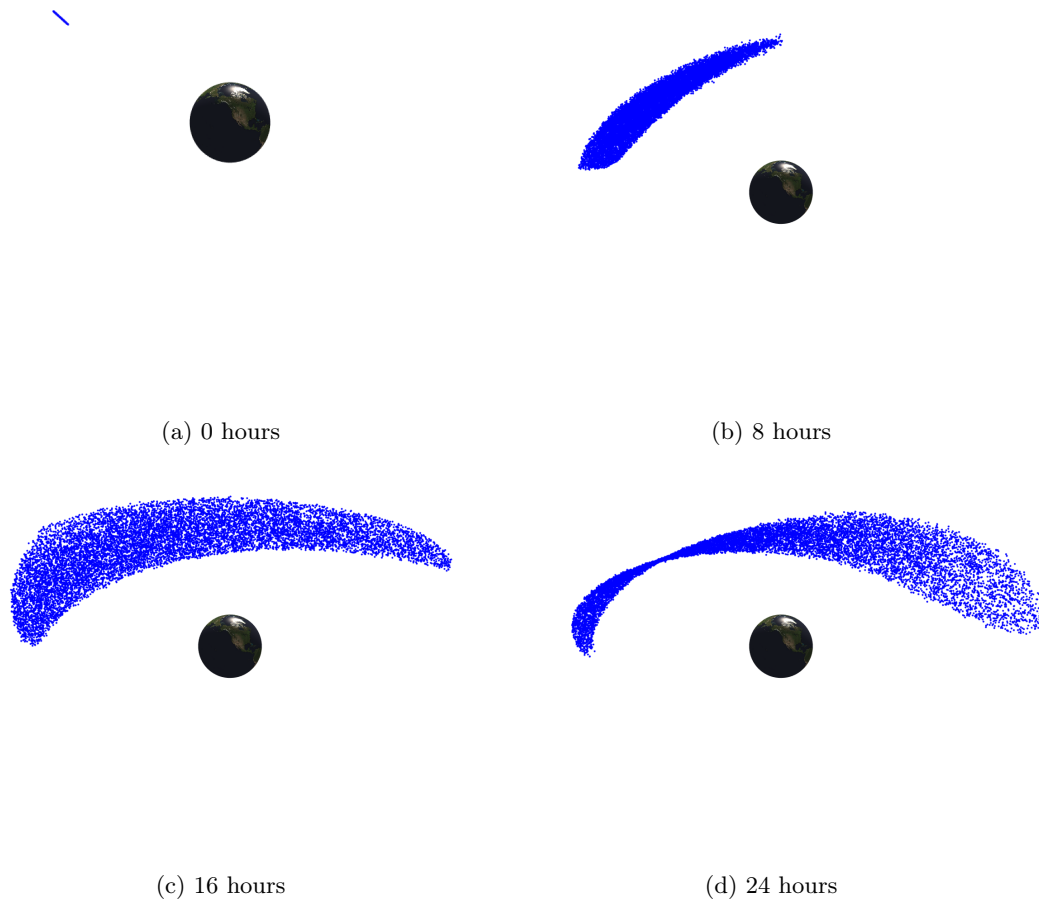


Fig. 3. Constrained IOD evolution. Sub-figure (a) shows the initial set of samples which is tightly constrained in angle, but not range. The propagation of these samples in sub-figures (b)–(d) show that the initial uncertainties in range and velocity (range rate and angles rates) need to be managed by further observations.

very long time to explore the whole state-space. On the other hand if the proposal is large compared to the extent of the target distribution then the acceptance rate will be very low reducing the diversity of the resulting set of samples. The added complication for the IOD problem is that there are typically high correlations in each dimension of the state-space, meaning that a Gaussian proposal distribution which is independent in each dimension will not be effective. We have found that the adaptive MCMC algorithm of Haario [12] works very well for this application. By starting with a small proposal and adapting to the target distribution, an efficient proposal is generated by the chain itself. This method is appealing because it generalises in a straight-forward way to other measurement models, such as radar, which might measure range and range-rate along with angles. An example of the proposal adapting to the target distribution is shown in Figure 2.

The evolution of position uncertainty from IOD performed on a near GEO RSO is shown in Fig. 3. The CAR apogee and perigee values were constrained to be 60,000 km and 35,000 km respectively. The large state space explored by the IOD sampler is clearly shown. The figure highlights the need for dynamic sensor control to manage sensor resources between newly formed RSOs and established RSO tracks.

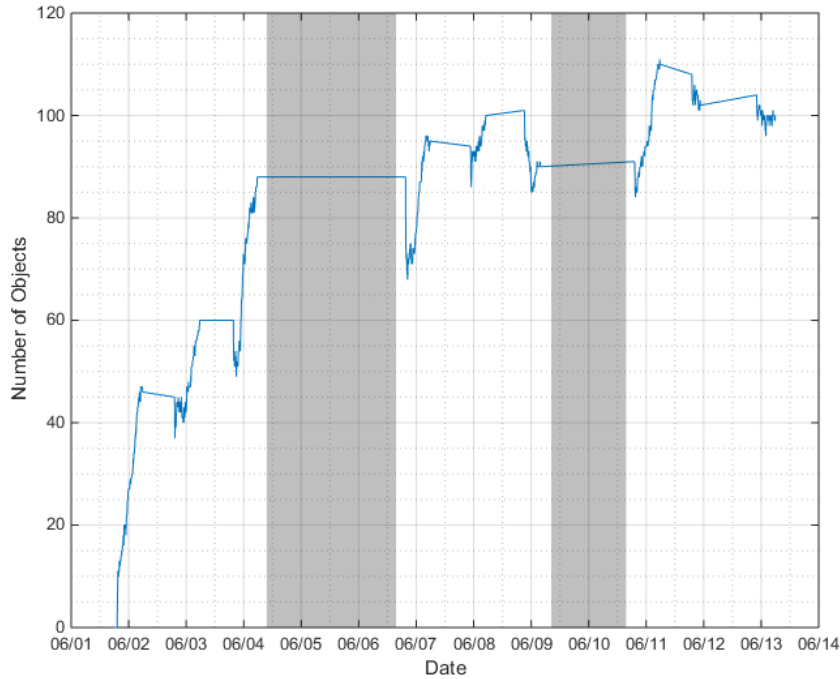


Fig. 4. The number of objects tracked during the ten night trial. The grey sections highlight nights where observations could not be made due to weather.

6 RESULTS AND DISCUSSION

On June 1st, 2017 a multiple day trial was conducted to assess TOTIS. TOTIS was tasked to conduct surveillance of the GEO belt between longitudes 85° and 105° and between latitudes $\pm 2.5^\circ$; apogee and perigee limits on objects were set between 60,000 km and 35,000 km respectively. TOTIS was initialised with an empty catalogue.

The multiple day trial was run over ten nights, however during two periods no observations were taken due to poor weather conditions. Without any human intervention TOTIS was able to construct and track a catalogue of over 100 RSOs from the surveillance region. The evolution in the number of RSOs tracked during the trial is shown in Fig. 4, time periods highlighted in grey indicated nights when TOTIS could not take observations. The location of RSOs relative to the surveillance regions at the end of the trial is shown in Fig. 5. The green circles are the centres of the segments of the sky TOTIS was tasked with surveying. The blue cross is the position the sensor was targeting at this point in time - an RSO update task, rather than a surveillance task. The groups of coloured dots are RSO positions mapped into measurement space. As this data was taken at the end of the night the uncertainty of the majority of RSOs is small.

Throughout this trial TOTIS tracked multiple RSOs that were not geosynchronous objects with an orbital period less than 24 hours. These objects gradually moved outside of the predefined surveillance region. As seen in Fig. 5, TOTIS was able to follow them outside the region and reduce their uncertainty as necessary. At the end of the trial TOTIS, while tracking over 100 RSOs, still had spare capacity to survey regions of the sky for potential new RSOs.

Figure 6 shows the standard deviation of the position error reported by TOTIS for three typical objects that TOTIS tracked during the trial. Sharp vertical decreases in uncertainty represent when TOTIS observed that particular RSO. The periodic segments with positive slope represent growth in uncertainty during daytime when the sensor is unable to take observations.

During the daytime periods between observation nights the objects can not be viewed, hence the large increase in state uncertainty between these times. However, the RSOs were still able to be reacquired by TOTIS the following night demonstrating its ability to perform continual surveillance without human




Fig. 5. Tracked RSOs at the end of the ten night trial. Coloured groups of points represent the state of the RSOs in measurement coordinates. The blue star is where the telescope has been tasked to look. The green circles show the directions that the telescope could point when tasked for surveillance.




Fig. 6. The standard deviation of position error reported by TOTIS for three selected objects. The grey sections highlight nights where observations could not be made due to weather.

Table 1. The standard deviation of position error reported by TOTIS for three selected objects at the end of the trial.

	RSO₁	RSO₂	RSO₃
Std. Position_x^{eci} (m)	220.1	257.5	603.8
Std. Position_y^{eci} (m)	1611.4	700.1	5273.7
Std. Position_z^{eci} (m)	248.2	152.3	134.2

intervention. It should be noted that a few RSOs could not be reacquired and were automatically deleted from the catalogue.

The standard deviation of position error reported by TOTIS at the end of the trial period are shown in Table 1, for the three selected RSOs. At this point in the trial the uncertainty of the RSOs is relatively stable. As an orbital slot at geostationary is approximately 73 kilometres across this uncertainty is adequate to make reasonable assessments about the location of RSOs. If further refinement is needed the sensor can be directly tasked for this purpose. The radial direction of standard deviation has the largest component in the y inertial axis at this point in time. An optical sensor is unable to directly measure the radial component of the state, hence y has a larger value relative to the other components.

TOTIS was able to find and continue tracking RSOs after multiple days without viewing opportunities. The particle filter’s ability to more precisely capture the uncertainty in the orbits allows TOTIS to accurately predict the positions, and then reacquire, RSOs after long periods without observations.

The catalogue constructed by TOTIS was compared to the Two-Line Element catalogue published by NORAD. There was a clear correspondence between TOTIS’s catalogue and NORAD’s throughout the duration of the experimental campaign. By comparison with the NORAD catalogue it was found that the RSOs selected for analysis above can be identified as RSO₁: ChinaSat 10 (32478), RSO₂: Express AM33 (37677), and RSO₃: Intelsat 15 (36106).

7 CONCLUSION

With space situational awareness becoming an ever more demanding and important task, developing systems that can autonomously produce catalogues is critical. TOTIS has demonstrated the ability to effectively and efficiently construct, refine and manage a catalogue of near-geosynchronous objects over a two week period. TOTIS employs state estimation techniques that allow for uncertainty of an object’s state to be captured with greater accuracy and sensor control algorithms to efficiently task SSA sensors that maximises the use of sensor resources.

REFERENCES

- [1] T. Hobson, L. Clarkson, T. Bessell, M. Rutten, N. Gordon, N. Moretti, and B. Morreale, “Catalogue creation for space situational awareness with optical sensors,” in *Proc. the Advanced Maui Optical and Space Surveillance Technologies Conf. (AMOS)*, 2016.
- [2] T. Hobson, N. Gordon, M. Rutten, and T. Bessell, “Dynamic steering for improved sensor autonomy and catalogue maintenance,” in *Proc. the Advanced Maui Optical and Space Surveillance Technologies Conf. (AMOS)*, 2015.
- [3] M. Rutten, J. Williams, N. Gordon, J. Stauch, J. Baldwin, and M. Jah, “A comparison of JPDA and belief propagation for data association in SSA,” in *Proc. the Advanced Maui Optical and Space Surveillance Technologies Conf. (AMOS)*, 2014.
- [4] N. Moretti, M. Rutten, T. Bessell, and B. Morreale, “Automated resident space object catalogue construction and maintenance using optical sensor management,” in *Proc. International Astronautical Congress*, 2017.

- [5] S. Särkkä, *Bayesian Filtering and Smoothing*. Institute of Mathematical Statistics Textbooks, Cambridge University Press, 2013.
- [6] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, “A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking,” *IEEE Transactions on Signal Processing*, vol. 50, no. 2, 2002.
- [7] A. Doucet, S. Godsill, and C. Andrieu, “On sequential Monte Carlo sampling methods for Bayesian filtering,” *Statistics and computing*, vol. 10, no. 3, pp. 197–208, 2000.
- [8] C. Musso, N. Oudjane, and F. Le Gland, “Improving regularised particle filters,” in *Sequential Monte Carlo methods in practice*, pp. 247–271, Springer, 2001.
- [9] D. Gamerman and H. Lopes, *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference, Second Edition*. Chapman & Hall/CRC Texts in Statistical Science, Taylor & Francis, 2006.
- [10] M. D. Schneider, “Bayesian linking of geosynchronous orbital debris tracks as seen by the Large Synoptic Survey Telescope,” *Advances in Space Research*, vol. 49, no. 4, pp. 655–666, 2012.
- [11] K. J. DeMars, M. K. Jah, and P. W. Schumacher, “Initial orbit determination using short-arc angle and angle rate data,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 3, pp. 2628–2637, 2012.
- [12] H. Haario, E. Saksman, and J. Tamminen, “An adaptive Metropolis algorithm,” *Bernoulli*, vol. 7, no. 2, pp. 223–242, 2001.