Optimal SSN tasking to enhance real-time Space Situational Awareness

John Ferreira III, Islam Hussein, and Joseph Gerber
Applied Defense Solutions
Robert Sivilli
United States Air Force Research Laboratory

ABSTRACT

Space Situational Awareness (SSA) is currently constrained by an overwhelming number of resident space objects (RSOs) that need to be tracked and the amount of data these observations produce. The Joint Centralized Autonomous Tasking System (JCATS) is an autonomous, net-centric tool that approaches these SSA concerns from an agile, information-based stance. Finite set statistics and stochastic optimization are used to maintain an RSO catalog and develop sensor tasking schedules based on operator configured, state information-gain metrics to determine observation priorities. This improves the efficiency of sensors to target objects as awareness changes and new information is needed, not at predefined frequencies solely. A net-centric, service-oriented architecture (SOA) allows for JCATS integration into existing SSA systems. Testing has shown operationally-relevant performance improvements and scalability across multiple types of scenarios and against current sensor tasking tools.

1. INTRODUCTION

The advancing technologies utilized by surveillance systems provide increasing opportunities to observe and track the over 500,000 resident space objects (RSOs) of 1 cm size or greater. This improvement in capability is immensely beneficial to Space Situational Awareness (SSA) but also introduces new challenges to the execution of SSA tasks. The two primary concerns are being able to maintain this number of objects simultaneously at a catalog level of accuracy, and the ability to process all the observations required to meet that objective in a timely fashion. Current methods of surveillance system tasking are dependent on past historical data, determining the number of observations required based on studies of historically achieved accuracies. Combined with infrequent tasking updates and inconsistent sensor response rates, this often causes objects to be under- or over-collected, which is an inefficient use of the limited resources. Furthermore, when the level of observations collected, already at hundreds of thousands a day, is increased by multiple factors the challenge of integrating them and maintaining an RSO catalog using the current, manually-intensive methods becomes overwhelming.

The Joint Centralized Autonomous Tasking System (JCATS) is a new and innovative tool that Applied Defense Solutions, Inc. (ADS), has developed to address these specific needs of the SSA community. Initially developed under an Air Force Research Laboratory (AFRL) Small Business Innovation Research (SBIR) project, JCATS implements new analysis tools and techniques to automate the SSA workflow as much as possible and improve the effective use of available resources. First, two new tools, leverage from previous AFRL efforts, were implemented as the core analysis engine for JCATS. A data fusion engine called the Finite Set Statistics (FISST) is used to take observations and maintain an RSO catalog. Due to the multi-hypothesis nature of FISST this allows for various orbit events, such as maneuvers, to be analyzed and considered during observation processing. The Information State Receding Horizon Control (ISRHC) tool is used to optimize a surveillance system tasking schedule based on information gain. Secondly, a big innovation that JCATS has implemented is the use of information-based metrics for analysis. Instead of specifying that an object be observed at certain average frequency the observation schedule is determined by the requirements set, such as a threshold position uncertainty level or the proximity to another object. As these requirements are calculated they raise or lower the effective priority of an object, which in turn influences the amount of information gain from a set of observations. Thirdly, JCATS' approach to the SSA problem offers alternative standard operating procedures to maximize effectiveness. The transfer of data and instructions between components is primarily driven through a net-centric, service-oriented architecture (SOA). This allows for a design that can provide rapid response capabilities that adapt to real-time circumstances.

2. FILTERING AND UNCERTAINTY PREDICTION

The use of an information-based, centralized scheduling solution enables a fundamental change to the way observations are tasked for collection and the information extracted in order to achieve higher fidelity SSA. JCATS leverages recent AFRL work on FISST [1], a heterogeneous data fusion engine that dynamically updates an RSO catalog based on raw measurement. JCATS also uses ISRHC to find an optimal solution which maximizes the information content of the collected observations [2]. ISRHC requires an objective function to optimize. JCATS uses a FISST-based information entropy metric that seeks to meet user specified accuracy levels on each RSO. Once the accuracy level is met for a target, no SSN resources would be allocated to it, hence freeing up resources for meeting the requirements on other targets. More specifically, based on user requirements, each RSO with ID number i is assigned a desired level of tracking accuracy σ_d^i. This in turn results in a *desired* covariance matrix $\Sigma_d^i$for each RSO, where Σ_d^i is a 3 by 3 diagonal matrix with the diagonal elements being σ_d^i. More complex scenarios can also be accommodated, for example allowing the user to define in-track accuracy levels different from cross-track ones. For RSOs with no user specified accuracy requirements, the system automatically assigns them a default desired accuracy level to ensure that their uncertainty does not grow too large to the point that the RSO gets lost. This will not necessarily prevent any given RSO from getting lost. This is due to the fact that RSOs with higher priority may result in taking away SSN resources from lower priority level objects with a large uncertainty.

The FISST inference component has two main tasks. The first one is to use incoming raw sensor data, along with sensor meta-data such as biases and error covariances, to update an internal statistical database of the RSO population being tracked. In its most basic form, this database stores the *information state* of an RSO, which is composed of its mean position/velocity along with the associated covariance matrix. As a Bayesian technique, FISST is continuously executing two basic operations as shown in Fig. 1. The first is to update the internal database at observation collection times. This step usually results in a reduction in uncertainty in the observed object's state. In between observations, the internal statistical database is propagated (i.e., uncertainty propagation) according to established astrodynamic propagation models. This step usually results in increased uncertainty in the RSO information state.

FISST is, essentially, a rigorous hypothesis management technique. Hypotheses it accounts for include: object birth and death, data association hypotheses, misdetection and false alarm hypotheses, as well as object classification hypotheses. Given a hypothesis, FISST uses any of the existing uncertainty propagators, such as the Unscented Kalman Filter (UKF) or the Particle Filter (PF) to propagate and update the underlying RSO information states for that hypothesis. JCATS currently is capable of using either the UKF or the PF, or an intelligent combination of the two for maximal computational efficiency, to process the hypotheses. Unlike ad hoc techniques such as Multiple Hypothesis Tracking (MHT), FISST is mathematically rigorous which enables it to take into account that all such hypotheses are statistically dependent. FISST is capable of using such dependencies to further assert or reject hypotheses based on incoming observations. This allows FISST to extract the most amount of information from the data.

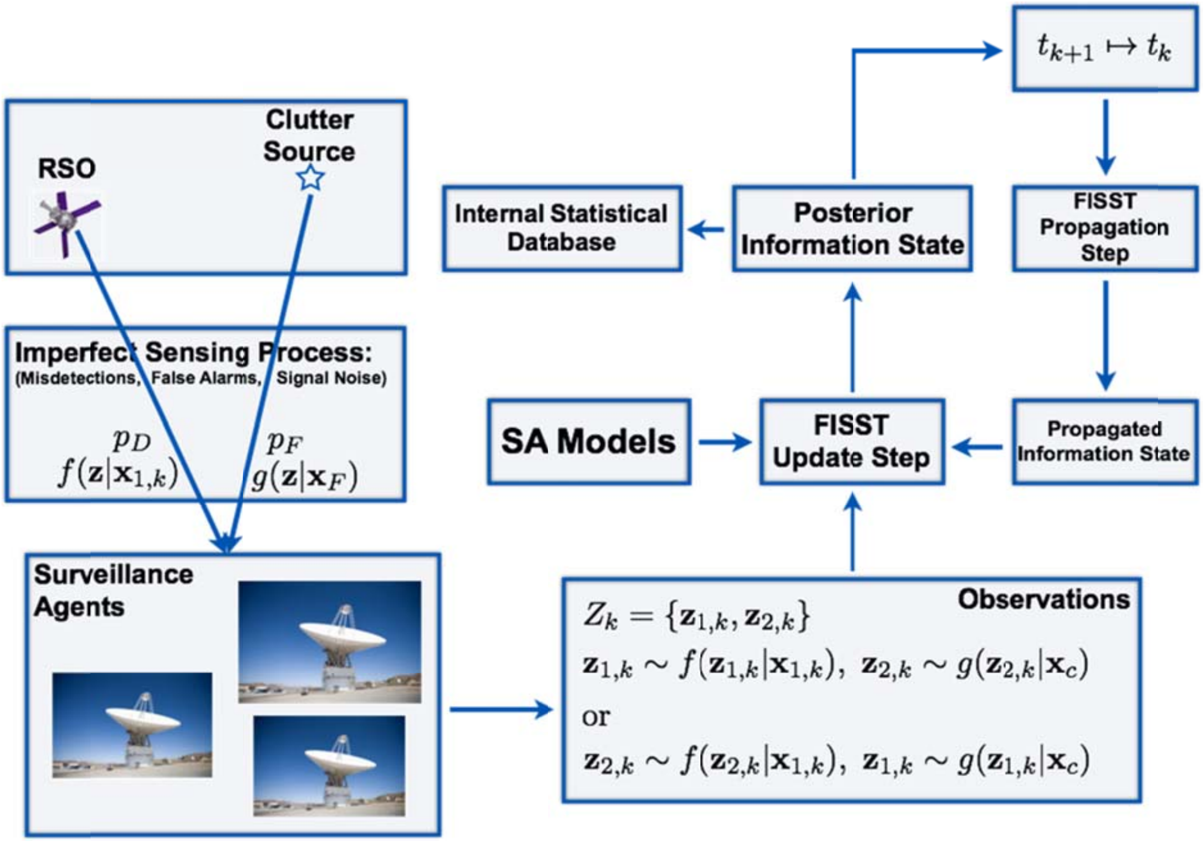

Fig. 1. FISST's Internal Information Flow

This informational efficiency inherent in FISST comes at the cost of CPU time efficiency. FISST-based children techniques such as the Probability Hypothesis Density (PHD) and the Cardinalized PHD (CPHD) [1], and to a lesser extent the δ-GLMB [3], involve heavy moment approximations to FISST to make it tractable. This however results in significant loss in FISST's informational efficiency discussed above. Research carried out by the authors and their collaborators in Academia and National Laboratories over the last 5 years has focused on introducing small approximations, mainly in the data association step, to FISST that enables it to be computationally tractable while retaining its informational advantages over alternative techniques [4]. Scalability of the resulting FISST algorithm is achieved by the randomization of the observation-to-RSO association problem within the update step. Reference [4] describes this approach in greater detail. The core idea behind the scalability of the proposed solution is that the discrete hypotheses are filtered according to an intelligent implementation of Markov Chain Monte Carlo (MCMC) sampling, called Smart Sampling MCMC (SSMCMC), that was developed jointly with Texas A&M University and demonstrated in [4]. While SSMCMC is an approximation technique, it results in a far less loss of information than those resulting from the severe approximations involved in the PHD and CPHD methods. This theoretical statement regarding information retention is currently being tested and verified in current research. The resulting algorithm is what is implemented in JCATS. This results in a scalable, real-time and highly informationally efficient inference engine within JCATS which provides significantly improved capabilities over the systems currently supporting the SSN mission.

The second task for the FISST component is to predict uncertainty for the purposes of scheduling. Hence, in the prediction step, one can account for object birth and death, false alarms and misdetections, and so on. During this prediction step, given a candidate multi-sensor schedule, JCATS simulates collected data in a way that conforms with the collecting sensor error statistics and properties. Hence, it is a realistic assessment of uncertainty performance over the scheduling window. We note here that an important difference between this scheduling prediction and raw data ingestion for online filtering is that in the former there is no data association to be solved. Since this is a simulation, and while JCATS does account for all the phenomena (sensor errors, object birth and death, etc.), the true data association is known. This helps reduce the computational efficiency of the scheduling step significantly.

3. OPTIMAL SENSOR SCHEDULING

ISRHC is the optimization component that seeks to find the optimal schedule over a future window of time. It was developed jointly by Texas A&M University and the Space Vehicles Directorate at AFRL in 2013 [2]. Reference [5] was the first to implement this in a sensor scheduling context. It is a stochastic gradient descent technique that seeks to minimize a given information-based cost function over a future time horizon. A stochastic gradient descent technique randomly selects a sensor tasking/schedule that decreases the information metric being optimized. ISRHC is fast since, unlike other stochastic search algorithms, it makes use of the "gradient" of the cost function. Thus, it requires fewer samples of the space of candidate solutions than non-gradient based methods as in, say, simulated annealing or neural network approaches. However, the optimal scheduling problem is a discrete optimization problem since the search space is composed of non-continuous discrete schedules and the mathematical space of schedules cannot be expressed as a smooth function of continuous parameters. Therein lies the power of ISRHC as it is able to probabilistically infer the "discrete gradient" as it explores the mathematical space of candidate solutions. This implies that ISRHC is able to find an optimal schedule rapidly.

Next, we discuss the information-based objective function that is minimized in JCATS. As previously mentioned, the FISST inference engine propagates and updates RSO covariances. Each RSO then has an *actual* covariance matrix Σ_a^i that is constantly being processed within the FISST component. The difference between the desired and actual covariances can then be used as a measure of information uncertainty deviation from a desired value. How to measure this deviation is a key question in information theory. Unlike in vector calculus, in matrix algebra one does not simply take the difference Σ_d^i - Σ_a^i between the two matrices as that does not necessarily produce a proper measure of uncertainty error. To produce a proper measure of uncertainty error one turns to the notion of information entropy, which is firmly founded in information theory. Any individual covariance has associated with it an entropy value which is a measure of how uncertain the system is of an RSO's state; the larger the entropy the more uncertain one is of the object's state. The desired and actual entropies are computed according to the standard formulae [6]:

$$E_d^i = \frac{1}{2}\ln((2\pi e)^n \left|\Sigma_d^i\right|$$

$$E_a^i = \frac{1}{2}\ln((2\pi e)^n \left|\Sigma_a^i\right|$$

The entropic error is then given by:

$$e^i = \begin{cases} E_a^i - E_d^i & if\ E_a^i > E_d^i \\ 0 & otherwise \end{cases}$$

Thus, whenever the desired entropy level (equivalently, the desired accuracy level) is not met the error is positive and otherwise it is zero. This can then be used as a proper measure of error between desired uncertainty and the actual uncertainty. This is the individual RSO entropic error given at any given point in time. The objective function implemented in this work is a weighted average of all the entropic errors for all RSOs that we call the *global weight entropic error (GWEE).* The weights are determined by a prioritization, or "category" system, defined by the user. JCATS can accommodate any category/prioritization system.

ISRHC seeks to find the optimal sensor-to-RSO schedule that best meets user-desired tracking accuracies by minimizing the predicted GWEE over a given future scheduling window of time. The GWEE can be predicted by employing the FISST engine to predict the RSO covariances over the scheduling window. Since FISST is being employed, the prediction is capable of taking into account phenomena such as the impact of misdetections (based on visibility and probability of detection models applied to future scheduling period conditions) and potential clutter and background noise, as well as the potential that RSOs are maneuverable. This allows for taking into account whether the sensor is being operated in a wide field-of-view or a narrow field-of-view mode and the impact these modes of operations have on improvement in RSO certainty.

JCATS can optimize over any desired performance criterion and not just information-based metrics. For example, it is capable of optimizing the SSN to meet user-defined revisit rates, or any other desired scheduling criteria. Our focus in these developments has been on information-based metrics such as position uncertainty or proximity to other objects. However, JCATS is modular and versatile enough to accommodate any other scheduling objective functions (including mixtures of criteria, such as a mixture of both revisit rates and information content simultaneously).

4. SYSTEM CONNECTIVITY

A key aspect to JCATS' design is its interactions between components via a net-centric, SOA-based interface. A catalog of messages has been developed that allows for an agile and adaptable system of components, and a high level workflow of the data transferred between them is seen in Fig. 2. First, operators have two portals that provide control and oversight of JCATS, the Mission Management Requirements (MMR) Portal and the Key Performance Parameters/Report Viewer Portal (KPP). The MMR is the tool with which the different information-based metric requirements can be configured against RSOs. The KPP allows an operator to view JCATS performance metrics, catalog information, and other values via reports or graphs. Each of these portals has a set of pre-defined SOA messages that are used to send the operator configured data into JCATS. This system allows changes made by operators to be ingested quickly without interfering with any actions JCATS is currently performing, while being taken into account for all future actions. This allows JCATS to be dynamically adapted as an operator's needs change in real time.

Similarly, a set of messages exists for JCATS to interact with surveillance agents directly. This allows for tasking to be assigned, observations to be collected, and similar function. The most important feature of the JCATS/Surveillance System relationship is a tight feedback loop. Actions taken at the sensor site, such as stopping collections due to bad weather or missed observations, are reported back to JCATS via SOA messages. This allows JCATS to automatically recalculate sensor tasking schedules, updated with the new knowledge of its resource availability and catalog state. To leverage this loop, JCATS tasks surveillance systems at a finer scale than most current operations, describing specific windows with which to view an object. This automated process helps remove "man-in-the-loop" requirements that exist in current system, letting more observations be processed and achieving more productivity from surveillance systems.

An alternative method of interaction between JCATS and surveillance systems is through an external tool network, such as ARCADE. The use of SOA-based interfaces allows for easy integration into existing systems such as this. Furthermore, due to JCATS' modular, JAVA-based design it is also possible for external tools to call specific JCATS functionality directly, leveraging the power of FISST or ISRHC against existing catalogs.

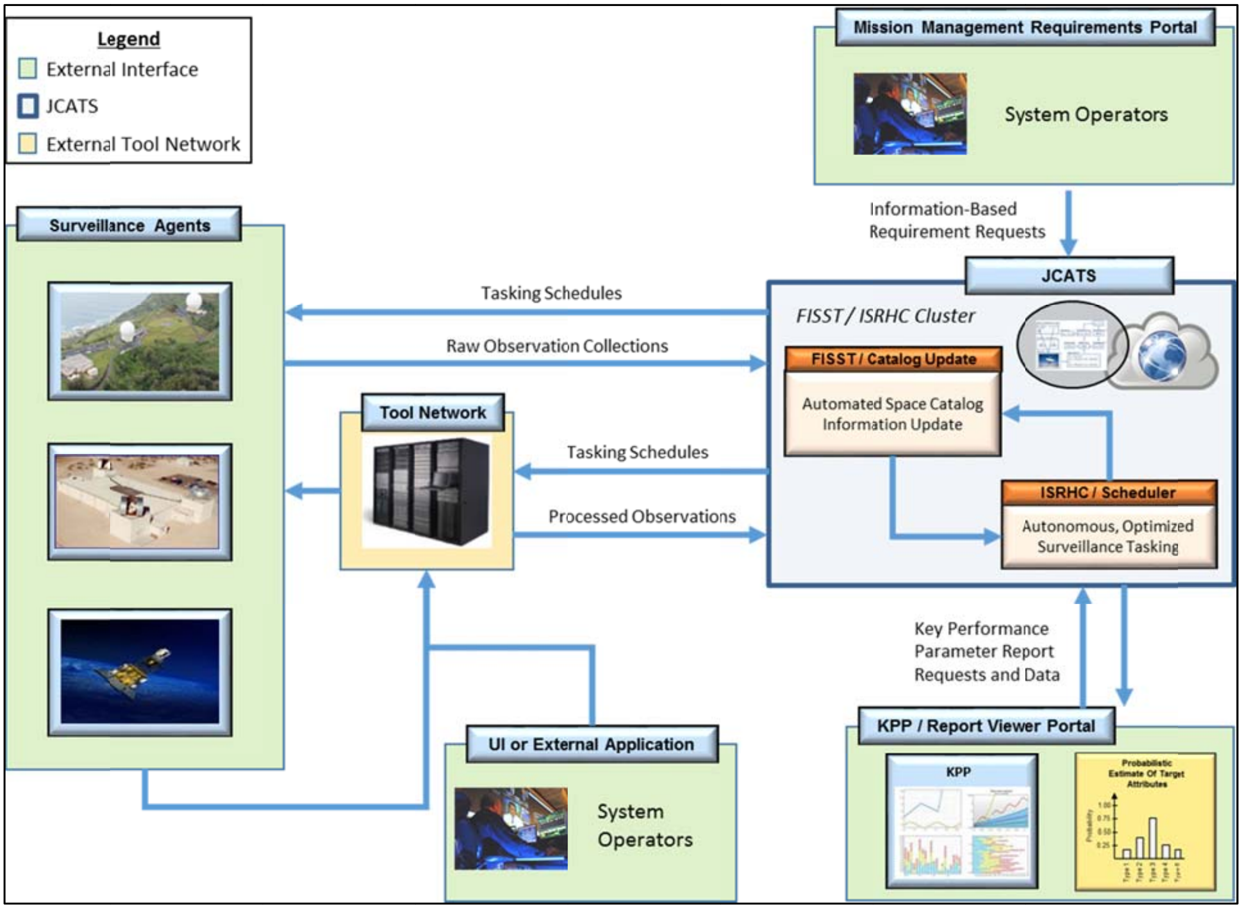

Fig. 2. Data Transfer Workflow

5. DEMONSTRATION RESULTS

In the SBIR effort JCATS was tested and analyzed in several different ways. First, the capabilities of the tool against the size of this combinatoric problem were demonstrated. The system is linearly scalable in terms of the three major inputs: number of RSOs, number of surveillance systems, and scheduling length. This was tested for catalogs that constituted thousands of objects. JCATS was also able to create schedules against these catalogs in operationally-relevant timeframes for schedules that were hours to days long. Fig. 3 shows a snapshot of the JCATS catalog, as it is compared to the "truth" values for all RSOs, for a test scenario. Objects that satisfy all requirements are colored green, with those that have become old going from yellow to red. The graph included shows the decrease of the average position error, from truth, for the entire catalog over time. This particular tested started with a new catalog and had to confirm/acquire each of the objects, causing the initial error growth.

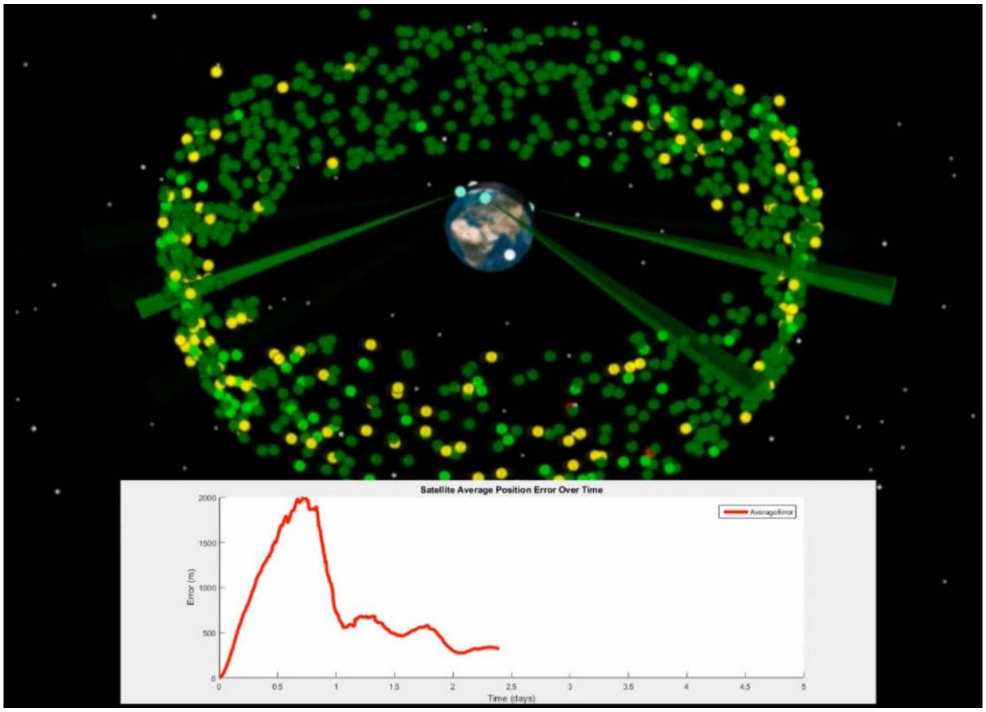

Fig. 3. A snapshot of the knowledge state JCATS has achieved while monitoring a new catalog of objects.

JCATS was also tested against multiple, realistic scenarios to demonstrate its scheduling capabilities against RSOs with dynamic requirements. Fig. 4 shows one of the result metrics, used to characterize JCATS abilities, for a single scenario. This shows the time between observations for a high interest object (HIO) when it is analyzed against a catalog of other objects with lower priorities. The behavior matches the anticipated results perfectly. Due to the scenario's geometry the HIO had short passes through several surveillance agents' fields of view with gaps between them. The 45-minute cyclic pattern shows this behavior, as well as the near constant target when in view. This scenario also demonstrated its ability to maintain the catalog of objects, with all other objects that could be observed maintaining roughly equal orbit uncertainty values by the end of the analysis period. These types of metrics were also used to compare JCATS to an existing sensor tasking tool, a next-generation version of one currently used in operations, and was able to demonstrate improvements in the final schedule, such as observing more objects and meeting more catalog requirements at the end of a given time period.

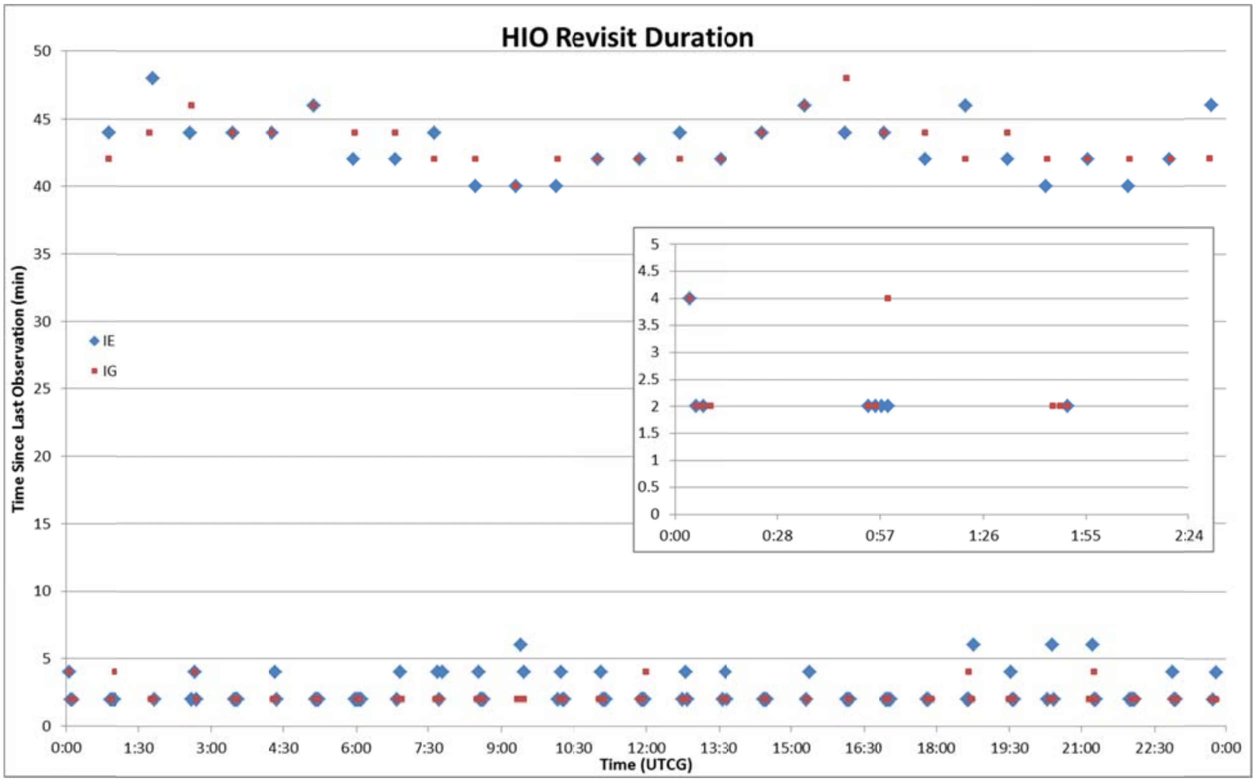

Fig. 4. The duration between observations of the HIO object for a single scenario.

6. CONCLUSION

JCATS provides a new tool, and a new perspective, on sensor tasking to solve the upcoming Space Situational Awareness problems. It does this by leveraging new analysis tools and techniques, framing the solution in terms of information-based metrics instead of coarse observation counts, and provides automated capabilities that can replace manually-intensive current procedures. The SBIR effort demonstrated the efficacy of JCATS across multiple types of scenarios, including catalog maintenance, and against existing sensor tasking tools. These advances will help enable the SSA community to expand its capabilities for observation collection and processing and, therefore, knowledge.

Currently JCATS is continuing to be used in several other efforts, expanding its feature set and to continue testing against catalogs. Future endeavors are being pursued to continue raising the technical readiness level of JCATS.

The authors would like to acknowledge support provided by AFRL SBIR Grant FA9453-15-M-0490.

7. REFERENCES

[1] R. P. S. Mahler, Statistical Multisource-Multitarget Information Fusion, Artech House, 2007.

[2] Z. Sunberg, S. Chakravorty and R. Erwin, "Information space receding horizon control," *Trans. on Syst. Man Cybernet. B,* 2013.

[3] H. G. Hoang, B. T. Vo and B. N. Vo, "A fast implementation of the generalized labeled multi-Bernoulli filter with joint prediction and update," in *Information Fusion*, Washington, D.C., 2015.

[4] W. Faber, S. Chakravorty and I. Hussein, "A randomized sampling based approach to multi-object tracking," in *IEEE Conference on Information Fusion*, Washington DC, 2015.

[5] I. I. Hussein, S. Chakravorty, Z. Sunberg, R. S. Erwin and M. K. Jah, "Stochastic Optimization for Sensor Allocation using AEGIS-FISST," in *AAS SFM Meeting*, 2014.

[6] T. M. Cover and J. A. Thomas, Elements of Information Theory, Wiley-Interscience, 1991.

[7] I. R. Goodman, R. P. S. Mahler and H. T. Nguyen, Mathematics of Data Fusion, Kluwer Academic Publishers, 1997.

[8] R. P. S. Mahler, Statistical Multisource-Multitarget Information Fusion, Artech House , 2007.

[9] W. Faber, S. Chakravorty and I. I. Hussein, "A Randomized Sampling based Approach to Multitarget Tracking," in *American Control Conference*, 2014.