

The Impact of Orbit Accuracy-Based Tasking on Sensor Network Efficiency

Neil K Dhingra, Cameron DeJac, Josh Neel, Alex Herz

Orbit Logic Incorporated, 7852 Walker Dr, Suite 400, Greenbelt, MD 20770

Trevor Wolf, Brandon Jones

The University of Texas at Austin, Austin, TX 78712

ABSTRACT

To support high-quality time-sensitive tracking of space objects, Orbit Logic's Heimdall Space Domain Awareness (SDA) sensor tasking software now supports planning for requirements on the expected track covariance after data fusion is performed. Operators can specify required levels of expected orbit accuracy for space object tracks at user-specified times and Heimdall will create sensor plans to achieve them, facilitating tactical operations or other time-sensitive missions requiring high orbit accuracy. Orbit accuracy is quantified as a user-specified function (e.g., the determinant) of the expected posterior track error covariance matrix obtained when the current track and planned sensor collects are combined using a specified filter (e.g., the Extended Kalman Filter). Heimdall employs a module developed by the University of Texas at Austin to perform these calculations and enable optimized planning around it. Crucially, intelligent planning for these requirements can reduce the amount of sensor time needed to maintain high quality tracks on space objects by ensuring that planned collects are scheduled at times and with parameters to generate a sufficient Value of Information (VOI) as measured in terms of orbit accuracy, and to prevent superfluous low-VOI tasking. This can be leveraged to efficiently achieve high-accuracy tracks on high-value targets and to efficiently reduce sensing on low-priority objects so that sufficient accuracy is maintained with fewer collects. To measure and illustrate the impact of this intelligent tasking, we will present studies using sample scenarios with relevant objects and representative sensor networks.

1. INTRODUCTION

The proliferation of space objects and the increasingly contested nature of space as a warfighting domain means that high-quality, often time-sensitive, tracks must be maintained on many objects to enable tactical operations and other missions as well as for effective Space Domain Awareness (SDA). Such elevated requirements impose a larger burden on sensor networks. However, intelligent planning can reduce the amount of sensor time required to maintain high quality tracks on space objects by scheduling sensor collects at times and with parameters to obtain specified levels of orbit accuracy, thereby reducing the burden on sensor networks and allowing them to effectively monitor more objects. In this paper, we will measure and illustrate the impact of orbit accuracy-based intelligent sensor tasking for SDA.

Orbit Logic's Heimdall SDA tasking software has been updated to support specified levels of expected space object orbit accuracy as a requirement for the generated sensor schedule. This can be leveraged to achieve high-accuracy time-sensitive tracks on high-value targets more efficiently, facilitating time-sensitive operations requiring high orbit accuracy. In addition, it can be used to intelligently reduce sensing on low-priority objects so that sufficient accuracy is maintained with fewer collects. Orbit accuracy is determined by the existing track and the parameters of the sensor collects on them, including timing, sensor quality, sensor phenomenology, viewing geometry, atmospheric conditions, other space or astronomical objects in the field of view, and others, and environmental conditions. By considering these factors when planning, sensor time can be preserved by scheduling data collection tasks that are expected to provide high VOI and omitting tasks that are expected to provide low Value of Information (VOI).

Rather than greedily maximizing the VOI for each individual collect, Heimdall's intelligent planning optimizes the schedule to support broad mission objectives. For example, it may schedule that a lower-but-sufficient VOI collect on one object that is relatively well-tracked to enable a higher-VOI collect on another that has a worse track. For high-interest space objects which may unexpectedly maneuver and cause large changes to their orbit, orbit accuracy requirements can be combined with persistent monitoring requirements to plan for high-quality tracks and prevent loss of custody with fewer sensing resources.

Specifications on orbit accuracy and other planning factors are expressed in sets of orders, tasking requirements which Heimdall performs intelligent sensor network tasking to fulfill. Each order specifies requirements on data gathered on a space object, identified by NORAD ID or a provided ephemeris. These requirements may entail observation timing within specified time windows, observation timing around orbital events, recurring observations at a given cadence, permissible sensor phenomenologies, permissible sensors, permissible viewing geometry, permissible weather during the collect, permissible viewing conditions (e.g., object apparent visual magnitude), minimum clearing radius, specified collection duration, minimum radar cross section, other permissible sensor parameters, and, now, a target orbit accuracy at given time(s). These orders can be overlapping and may each pertain to different aspects of the collected data. For example, an operator may issue an order requiring that data is collected at a given cadence and an order requiring that orbit accuracy meets a given threshold; Heimdall will plan accordingly to ensure that both requirements are met.

Alternatively, instead of in strict constraints, these factors can be used in components in the configurable SDA-specific figure of merit (FOM) to reflect catalog and/or mission objectives. Heimdall creates schedules to optimize this FOM while obeying order constraints. Heimdall can run several optimization algorithms in parallel to generate different sensor schedules and choose the schedule that scores best in terms of the FOM.

Heimdall's enhanced orders, which now support requirements to achieve specified levels of orbit accuracy at specific times, is enabled by a module developed by University of Texas, Austin (UT Austin). Orbit accuracy is derived from the expected posterior space object track estimation error covariance, computed by the UT Austin module. The software computes the matrix specifying this covariance using the current track, planned future sensor collects, and other conditions based on a user-specifiable filter; the defaults are the Extended Kalman Filter (EKF) and an epoch state filter that directly maps information gained to expected accuracy at the prescribed time. Orbit accuracy is determined as a user-specifiable function of the covariance matrix can correspond to physical properties of the covariance ellipse, such as the volume or major axis length.

In this paper, we will demonstrate that Heimdall's intelligent tasking enhances sensor network efficiency and results in plans that achieve the same or better orbit accuracy with the same or fewer sensor collects. The plans and results will be created and evaluated with model sensor networks representative of government and partner assets.

Outline

In Section 2, we discuss background of the SSA/SDA problem and the challenges with which this paper is concerned. Section 3 provides an overview of the Heimdall sensor scheduling software used to perform the studies in this paper. In Section 4, we provide an illustrative toy example to demonstrate the value of intelligent track accuracy-based tasking. In Section 5, we describe the assumptions, problem data, how the study was performed, and the results of the study. Finally, Section 6 provides concluding remarks and Section 7 contains references.

2. CHALLENGES IN SSA/SDA DATA COLLECTION

Effective SDA requires data collection that balances the monitoring of multiple heterogeneous objects by multiple heterogeneous sensors to satisfy multiple related objectives. Broadly speaking, the successful surveillance of any given object requires that tasking result in observations that are 1) of sufficient quality and 2) of sufficient quantity and density/regularity. High-quality data on objects are required so that they can be filtered for high-accuracy orbit determination and/or to provide alternate data products, such as those related to object characterization and conjunction assessment. In addition, data must be collected at a sufficient cadence to prevent one from losing track of the object. In this paper, we focus on the challenge of collecting data on routine objects that may be assumed to not be maneuvering aggressively. Accordingly, we do not consider the problem of associating observations to objects nor do we consider how tasking may change in response to a detected maneuver.

Requirements on data quality and quantity are coupled with one another, with the object, and with the sensors collecting data. Observations with better sensors will result in better data products and better tracks and, for objects less prone to maneuvers, may mean that less frequent observations are required to maintain a high-quality track. Tasking requirements can change with time; higher quality tracks may be required near potential conjunctions. Moreover, the value of information provided by a particular sensor is also dependent on factors such as the timing of data collection and existing knowledge about the object because the orientation and size of the sensor noise covariance ellipse in relation to the estimated covariance of the state estimate changes with the sensor mode,

observation geometry, and other factors. Finally, effective tracking of space objects should be balanced with the search for new, previously unobserved space objects.

We focus on the challenge of achieving high quality tracks on an object with as few sensor observations as possible. As tasking challenges interact in more complex ways, the need for better planning and scheduling is more pronounced.

3. HEIMDALL OVERVIEW

The Heimdall solution is intended to support an operation staff as part of a wider workflow enabling Battle Management Command and Control (BMC2). It specifically occupies the functional role of optimizing sensor tasking across a large number of ground and space sensors to achieve overall SDA-related objectives. Heimdall interacts with other components of a wider architecture using machine-to-machine interfaces utilizing plug-ins that allow the specifics of those interfaces to be easily updated, or even to become compliant with completely different interoperability standards in different systems. This has already been demonstrated via installation of the capabilities in multiple customer's systems. The primary interface to Heimdall is a web interface, accessible via standard browsers, through which all of the core administrative and operational features can be accessed.

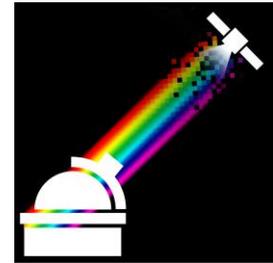


Figure 1: Heimdall Logo

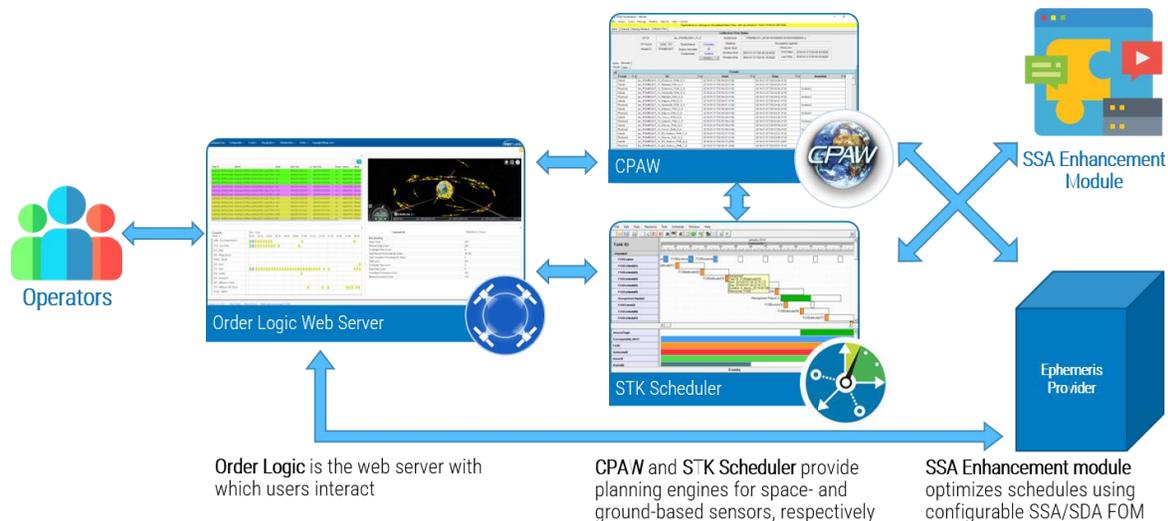


Figure 2: Heimdall System Architecture Diagram: Heimdall builds on and enhances mature Orbit Logic products to create optimized SSA/SDA sensor schedules

Scheduling/Tasking Algorithms

One of the core features of the Heimdall solution is the ability to generate coordinated, optimized observation schedules for the full set of available ground and space-based sensors for SDA observations. Ground sensor observation scheduling is performed by STK Scheduler scheduling algorithms, while space-based sensor observations are scheduled by Orbit Logic's Collection Planning and Analysis Workstation (CPAW) scheduling algorithms. Coordination between ground and space sensor planning is performed through process flow control by Heimdall the availability of observation fulfillment status through the shared Object Catalog database.

STK Scheduler provides multiple scheduling algorithms as well as an algorithm builder tool, to define refined algorithms for specific needs. In the SDA configuration, algorithms are fed the list of SDA FOM-scored observation opportunities and use that list as the basis for generating a high value, valid, deconflicted, coordinated observation schedule for all available ground sensors. Heimdall calls the STK Scheduler algorithms using an available STK Scheduler STK Connect command via its TCP/IP API with string keyword-value pairs. The specific

algorithm may be configured within Heimdall, but an option also exists to call an algorithm-builder-defined custom combination algorithm that computes solutions using multiple algorithms and returns the highest FOM-scoring solution. Earlier versions of the STK Scheduler algorithms were successfully demonstrated to CSpOC personnel as part of the SDA Software Suite from Analytical Graphics for a large scale SSN sensor tasking problem (10,000 objects, 24 hour schedule, 30 sensors), with optimized observation schedule solution time under 2 minutes.

CPAW, the component responsible for space sensor planning, has a similar set of algorithms for tasking schedule generation. Multiple algorithms are fed the SDA FOM-scored observation opportunities and iterated with high fidelity space sensor models to generate a high value, valid, deconflicted, coordinated observation schedule for all available space-based sensors. The nine available CPAW algorithms may be configured on or off via the Heimdall API, with the algorithm solution from the highest SDA FOM-scoring plan returned. CPAW scheduling algorithms are called via the available CPAW API using command strings delivered via TCP/IP interface. Scheduling results are saved directly to the Heimdall Object Catalog database, associated with applicable objects.

Heimdall controls the order of calls to CPAW and STK Scheduler for ground sensor and space sensor observation schedule generation, respectively, in order to create a coordinated observation plan across all available space and ground sensors. Alternatively, CPAW can be used to plan ground and space sensors together. Fulfillment status based on planned observations is stored in the Heimdall Object Catalog database to support this coordination. The CPAW and STK observation schedule generation algorithms are also available for optional use by Tasked and Contributing sensors for their own local sensor scheduling via a web interface.

SDA-specific Figure-Of-Merit

Heimdall makes use of an SDA-specific Figure-of-Merit (FOM). The SDA FOM scores each observation opportunity based on inputs (such as predicted information gain) from the Task Prioritization component and other factors (such as computed object visual magnitude), time since last observation, orbit covariance, anomalous behavior rating, and more.

Each factor has an associated configurable weighting attribute to specify the importance of the FOM factor relative to other FOM factors. Weighting attributes may be set to any value, including 0 (ignored) and negative (penalty) values, allowing for virtually unlimited tuning of the scoring FOM.

Additionally, the FOM is split into object factors and search area factors (as well as common factors that apply to both), and the scores for objects and searches are normalized against each other. Lastly, configurable weighting factors allow for the importance of object observations vs. searches for new objects to be defined.

The SDA FOM is tightly coupled within the SDA versions of STK Scheduler and CPAW. All observation opportunities are automatically scored using the configured SDA FOM as part of the standard processing flow in both software tools. The FOM for CPAW (space sensors) and STK Scheduler (ground sensors) are separate, allowing for different configuration/factor weighting values for each.

In a future version of the architecture the SDA-specific FOM will also be made available via web interface for optional use by Tasked and Contributing sensors for their own local schedule optimization.

Value of Information-Based Tasking

Incorporating measures of information gain into space-object sensor tasking procedures provides a way to quantify the quality of candidate observation opportunities. Heimdall was updated to enable tasking is informed by metrics related to the expected state error covariance of a space-object at a desired epoch time. This feature generates the expected state covariance matrix at that time provided an initial state covariance matrix and a set of candidate measurements. In addition to intelligent tasking, this feature provides elevated operator awareness of the expected catalog state and the tasking algorithm's rationale.

Minimizing the size of the covariance matrix corresponds to maximizing the information gained with a measurement sequence. The user, in Heimdall, will add a "Final Orbit Accuracy" to the order and the planning software will plan to achieve it. This parameter is by default the volume of the covariance matrix, but other metrics can easily be configured. Heimdall provides the initial state and state covariance matrix of a space-object, the observing asset type and location (both ground-based and space-based observing assets are acceptable). A candidate measurement

schedule is provided by the user which lists both a sequence of observation times, and the observing asset used per time.

Two renditions of the software were developed. The Extended Kalman filter (EKF) version sequentially updates the space-object's state covariance per measurement in the observation sequence. That is, for each candidate measurement in the sequence, the EKF algorithm evaluates a covariance matrix decrement, which is related to the expected information gained from said measurement. This decrement is then subtracted from the predicted covariance at the time of the measurement. The process is repeated for each measurement in the set. The TurboProp library is used to propagate the space-object state and state covariance between times in the measurement set. After all the measurements are processed, the state and state covariance matrix are propagated to a final epoch of interest, and this final covariance matrix is used to ensure the desired orbit accuracy is met [11]. The second version of the software uses an epoch-state filter formulation to calculate the final state covariance matrix. This formulation calculates a covariance matrix decrement in a batch-like formulation, forgoing the recursive procedure of the standard EKF.

A necessary component of the project was accurately including process noise in space-object dynamics modeling. Process noise is important in quantifying how much information is lost in propagating from measurement-to-measurement – or in other words, how much the covariance matrix grows between measurements. To do this, a new tool was developed in the TurboProp library for propagating a process noise transition matrix between candidate measurement times. This transition matrix was then incorporated into the standard EKF formulation and the epoch-state formulation of the software. The software was tested and validated on both ground-based and space-based sensor tasking scenarios.

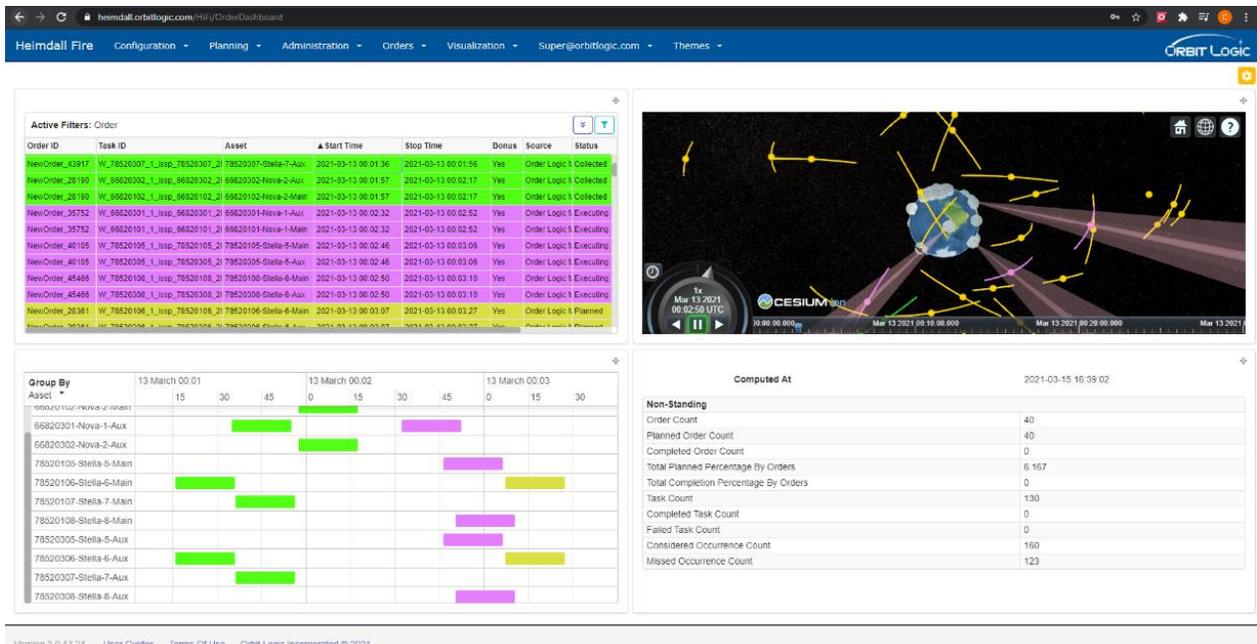


Figure 3: Heimdall Dashboard Page with the light color theme

External Plan Ingest

Heimdall was deployed and made available to external users via an Order Logic hosted machine configured for interfacing with leading commercial SSA operators (LeoLabs and Numerica, now Slingshot). Commercial SSA operator observation plans were retrieved and ingested by Heimdall, converted to ISSP format, and then utilized to show how commercial plans can inform Department of Defense (DoD) SSA sensor observation planning to meet DoD operational objectives, including meeting specific orbit accuracy goals.

Heimdall User Interfaces

Order Logic was developed as a user-facing interface for Orbit Logic's planning software application. The web application has previously been configured as the program-specific front end for both STK Scheduler and CPAW

planning applications (for ground and space-based assets, respectively). In Heimdall, Order Logic is configured to interface with both the STK Scheduler and CPAW planning engines, and has additionally been enhanced to provide overall workflow and automation control.

Providing an SDA-beneficial software automation framework for a distributed sensor network with worldwide non-traditional sites necessitated a web-enabled solution – one with the ability to monitor the state of space environment from many coordinated consoles and manage data flows in a highly configurable manner. As such, the web-based Graphical User Interfaces (GUIs) comprising the Heimdall solution are key to the overall operations concept.

One of the primary user features exposed through the web interface is visualization of the sensor tasking plans. Heimdall provides multiple ways for an operator to view, explore, and understand planned SDA tasking for ground and space sensors.

A configurable dashboard table view dynamically presents observations in time order, highlighting observations in progress (either in real-time and/or simulated time) and moving through the list of observations as time progresses. The presented list of observations can be filtered based on user preferences. The same Dashboard page provides a more global perspective in a 3D visualization pane. Driven by Cesium, this view is normally configured to run in real-time as a companion to the table view on the Dashboard, showing observations in an accurate graphical view as they occur throughout the collection of available sensors. The user may also select specific observations in the table view, and the Dashboard page Cesium 3D view automatically zooms in on the associated sensor resource and forwards to the time of the selected observation to display a static view of the specific observation geometry.

The Heimdall table and 3D views are driven by the latest object catalog database and associated planned observations saved within the object data there. The screenshot in Figure 3 shows the table view and associated configurable filter, along with the embedded 3D Cesium view and associated metrics.

Configuration Manager

The Configuration Manager component of Heimdall provides the ability for authorized users (administrators) to define and configure permissions for users, add and configure new SDA sensors, specify sensor downtime, specify optimization goals, review performance metrics, and perform other related setup and configuration functions. Changes made within Heimdall configuration pages are stored to the associated Heimdall database for use internally and/or used to send Application Programming Interface (API) configuration commands to some of Heimdall solution component applications.

Visibility Computations

At the start of the planning process, constrained access computations are performed for each valid sensor/object combination. Computations consider line-of-sight visibility, lighting constraints (when applicable), sensor capabilities, sensor field-of-regard, object attributes, and any applicable object/sensor assignments and preferences and constraints. Because access computations for each object are independent of the access computations for other objects, these computations can be performed in parallel on many cores in order to speed computation time for large object catalogs.

4. ILLUSTRATIVE TOY EXAMPLE

To explain the advantage gained from value of information-based tasking, we present a simple toy example. We consider a scenario in which two sensors are available: one is akin to a radar and has a pancake shaped sensor error covariance (with high accuracy measurement of range and lower

accuracy measurement of angle) and the other is akin to an optical telescope and has a cigarette shaped sensor error covariance (with high accuracy measurement of angle and no measurement of range). For simplicity, we normalize these covariances so that they both have the same volume; i.e., the determinant of the error covariances are the same. With a naïve tasking strategy that only seeks regularly spaced observations, only one sensor is tasked because it is the only one available on regularly spaced intervals. Using an optimized tasking strategy, observations occur with a mix of the two sensors.

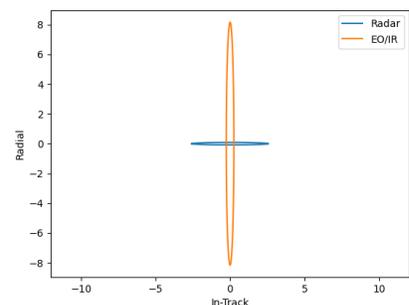


Figure 4: Measurement noise covariance for the radar-like and EO/IR-like sensor in the illustrative toy example.

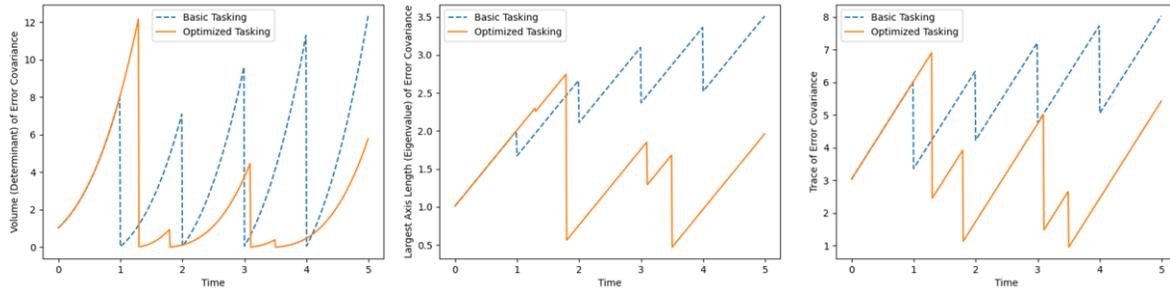


Figure 5: Effect of track metric-optimized sensor tasking.

The naïve tasking strategy with only radar-like observations results in state estimates that are much worse than the estimates that result from the optimized tasking plan. The estimation error covariance resulting from the naïve tasking strategy is larger by volume (i.e., the determinant of the estimation error covariance matrix), by the semimajor axis of the estimation error covariance (i.e., the maximum eigenvalue of the estimation error covariance matrix), and by the trace of the estimation error covariance matrix throughout the scenario with the sole exception being the interval after the naïve tasking tracker has ingested data and before the optimized tasking tracker has ingested data. By the end of the scenario, resulting estimation error covariance is 2.14 times larger by determinant, 1.79 by maximum eigenvalue, and 1.78 times larger by trace. Importantly, the covariance does not continue to grow when tasked intelligently, as the optimized tasking mixes sensing from diverse sensors to prevent covariance growth in particular directions, as is the case with the naïve tasking.

In addition, we consider how the confidence ellipse corresponding to the a posteriori estimation error covariance at the target time changes. We show the confidence ellipse in the plane of radial and in-track directions; by design, the cross-track direction in this toy example is identical to in-track for simplicity. Here, we consider observations occurring at the same time and compare the results of only using radar measurements versus alternating radar and EO/IR. We note that the results for only using EO/IR measurements are very similar.

Each of the radar measurements modify the estimation error covariance in similar ways, shrinking more in the radial direction than in the in-track and cross-track directions. This corresponds to the pancake shaped measurement noise from a radar wherein range is well determined and resolution normal to the line-of-sight to the target is limited. Of course, a sensor schedule with only EO/IR measurements would always primarily shrink the covariance in the in-track and cross-track directions, corresponding to the soda straw shaped covariance of the sensor. On the other hand, alternating measurements first shrink covariance primarily in the radial direction, then primarily in the in-track and cross-track directions, and so on. This alternation takes advantage of the particular strengths of each sensor and results in a better overall estimation error covariance; see the middle subfigure for a comparison of the uncertainty ellipses after identical numbers of measurements.

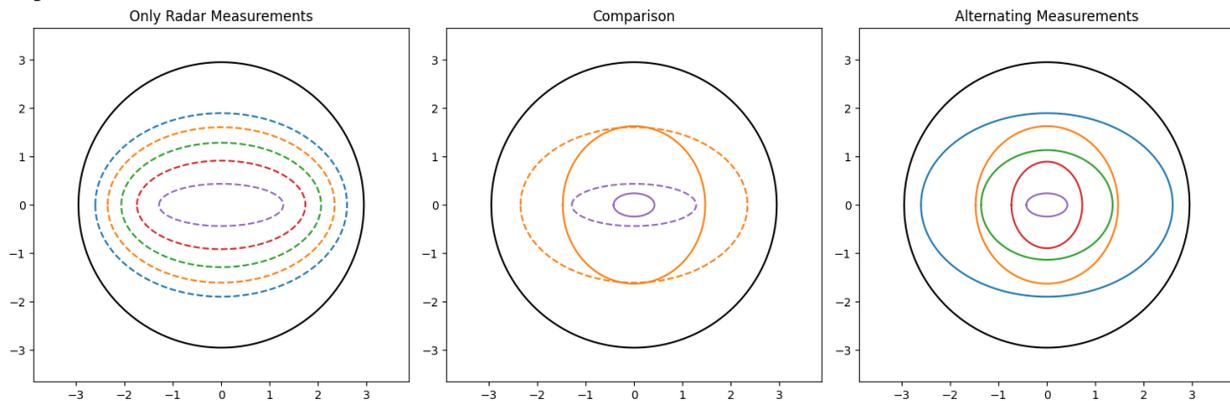


Figure 6: Expected a posteriori estimation error confidence ellipses for one standard deviation after fusing several observations. The black circle is the uncertainty ellipse with no measurements. The dashed lines trace confidence ellipses for the radar-only sensor schedule and the solid lines trace the confidence ellipse for the sensor schedule alternating radar and EO/IR.

5. STUDY DESCRIPTION

To provide real-world relevant results, we studied sensor tasking for a real object from the public catalog that has national security implications. In the interest of operational security, we are not disclosing the identify of the space object publicly. In addition, we considered tasking from three different assets based on models of SSN assets developed with publicly available data; similarly, we are not disclosing the identities of these assets in an abundance of caution. We employ the trace of the expected a posteriori estimation error covariance as the track accuracy metric and we plan to minimize it. Additional details are available upon request and similar studies for other objects, sensors, or larger missions are an area of interest for Orbit Logic. We encourage any interested parties to reach out to use about such continued work.

Better results for similar investment

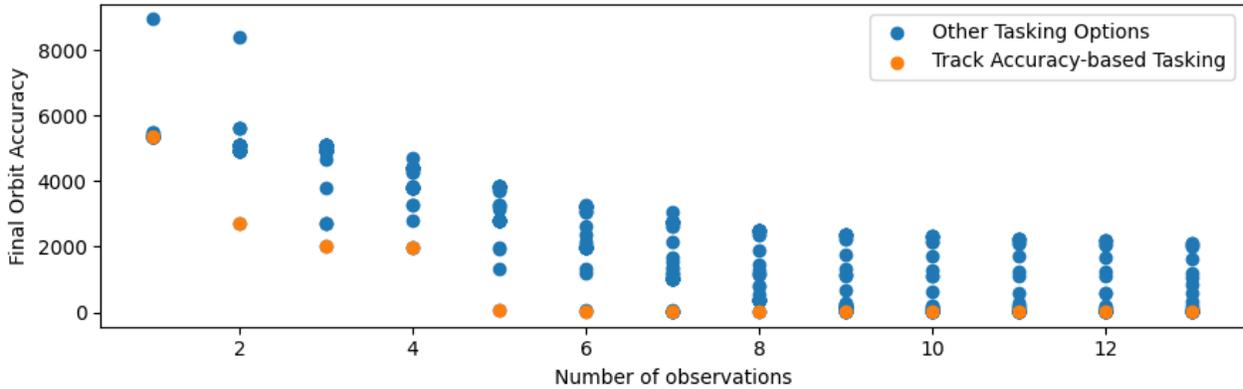


Figure 7: Track accuracy (trace of the a posteriori estimation error covariance matrix) at a target time for different potential sensor observation schedules for a specific RSO.

To compare the tradeoff between sensor time and track accuracy for different sensor schedules, we plot the expected track accuracy resulting from different plans. We randomly generated several representative schedules and compute the expected track accuracy resulting from them. In addition, we compute the expected track accuracy resulting from Heimdall's track accuracy-based tasking.

Figure 7 contains the results and demonstrates that track accuracy-based tasking forms a more efficient sensor schedule. For each number of sensor measurements, the track accuracy-based tasking achieves the best final orbit accuracy. This means that using track accuracy-based planning results in better bang-for-your-buck, i.e., better results for the same amount of sensor time devoted to an RSO.

Lower investment for similar results

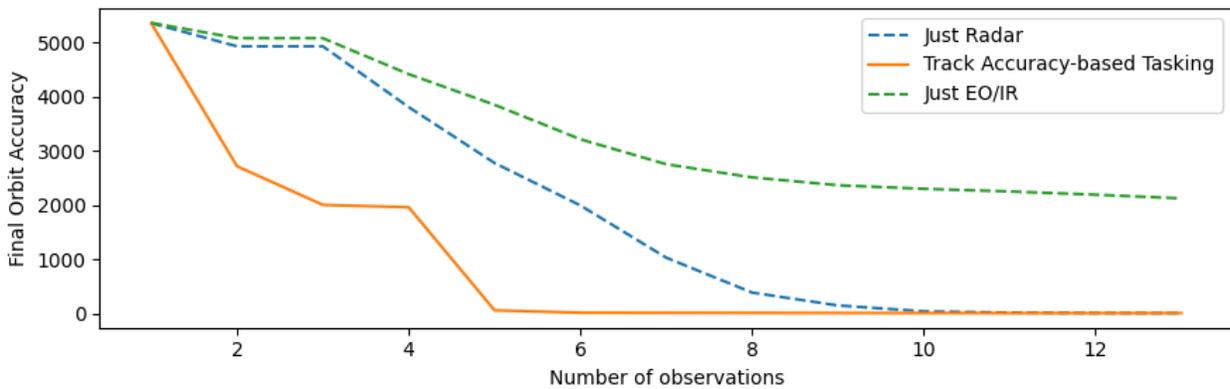


Figure 8: Track accuracy at a target time as more observations are fused. The observations' timing is identical.

The prequel also provides evidence that track accuracy-based tasking allows one to achieve similar levels of track accuracy for less sensor time investment. This allows operators to reduce the time devoted to an RSO with minimal deleterious consequences from track quality degradation.

To further illustrate this point, Figure 8 shows the track accuracy at a target time as more observations are fused. We compare to radar-only and EO/IR-only tasking. Clearly, track accuracy-based tasking reduces the sensor time needed.

EO/IR-only tasking results in a clearly worse final track accuracy. On the other hand, while radar-only tasking achieves high accuracy after many observations, track accuracy-based tasking achieves the same level of accuracy using much fewer observations. Moreover, track accuracy-based tasking provides a way to intelligently navigate the tradeoff between sensor time and track accuracy so that operators may reduce sensing resources used with less severe consequences for track accuracy. This may be useful to free up exquisite sensing assets for more sensitive missions while minimally impacting the catalog maintenance mission.

6. CONCLUDING REMARKS

Track accuracy-based tasking improves the efficiency of sensor networks performing SSA/SDA missions. It allows for better track accuracy with the same amount of sensing resources or, expressed as a slightly different but related requirement, allows one to achieve the same track accuracy with less sensing resources. These features provide clear advantages for SSA/SDA operators, allowing them to make better use of their sensor networks and better perform multiple simultaneous missions where sensor time is a scarce resource.

Finally, we note that we used number of observations as the measure of sensor resources for simplicity. Heimdall can account for weighting different assets in accordance with value, priority, money, etc.

Moreover, track accuracy is not the only component in the Heimdall FOM. Heimdall's FOM is configurable and contains many other factors that can drive planning, including the value of sensors, priorities, sensing geometry, predicted cloud cover, and many more. Heimdall creates sensor schedules to optimize this FOM, however the user has configured it; in this paper, we focused on the track accuracy component of the FOM for simplicity and to isolate the effect of intelligent planning on this specific factor.

7. REFERENCES

1. Alex Herz, Ella Herz, Kenneth Center, Doug George, Penina Axelrad, Shaylah Mutschler, Brandon Jones, "Utilizing novel non-traditional sensor tasking approaches to enhance the space situational awareness picture maintained by the Space Surveillance Network," Proceedings of the 2016 Advanced Maui Optical and Space Surveillance Technologies Conference, Maui, Hawaii, 2016.
2. Alex Herz, Brandon Jones, Ella Herz, Doug George, Penina Axelrad, Steve Gehly, "Heimdall System for MSSS Sensor Tasking," Proceedings of the 2015 Advanced Maui Optical and Space Surveillance Technologies Conference, Maui, Hawaii, 2015.
3. Blake, Travis, Michael Sanchez and Mark Bolden, "OrbitOutlook: Data-centric Competition Based Space Domain Awareness (SDA)," 30th Space Symposium, Colorado Springs, CO, 2016.
4. Donald B. Reid, "An Algorithm for Tracking Multiple Targets," *IEEE Transactions on Automatic Control*, AC-24(6):843-854, 1979.
5. Ronald P.S. Mahler, *Statistical Multisource-Multitarget Information Fusion*, Artech House, Boston, Massachusetts, 2007.
6. Brandon A. Jones, Steven Gehly, and Penina Axelrad, "Measurement-based birth model for a space object cardinalized probability hypothesis density filter," *AIAA/AAS Astrodynamics Specialist Conference*, San Diego, CA, August 2014.
7. Brandon A. Jones and Ba-Ngu Vo, "A Labeled Multi-Bernoulli Filter for Space Object Tracking," AAS 15-413, *AAS/AIAA Spaceflight Mechanics Meeting*, Williamsburg, VA, 2015.
8. Ba-Tuong Vo, Ba-Ngu Vo, and Antonio Cantoni, "Analytic implementations of the cardinalized probability hypothesis density filter," 55(7):3553-3567, *IEEE Transactions on Signal Processing*, 2007.

9. Steven Gehly, Brandon A. Jones, and Penina Axelrad, "An AEGIS-CPHD Filter to Maintain Custody of GEO Space Objects with Limited Tracking Data", *Proceedings of the 2014 Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui, Hawaii, 2014.
10. Gehly, S., B. Jones, and P. Axelrad, "A Sensor Allocation Scheme for Tracking Geosynchronous Space Objects," Accepted to *AIAA Journal of Guidance, Control, and Dynamics*, May 2016.
11. Hill, Keric, and Brandon A. Jones. "TurboProp Version 4.0." *Colorado Center for Astrodynamics Research*, 2009.
12. N. K. Dhingra, C. DeJac, A. Herz, T. Wolf, and B. Jones. Maximizing the Utility of Non-Traditional Sensor Network Data for SDA. In *Proceedings of the 2021 Advanced Maui Optical and Space Surveillance Technologies Conference*, Wailea, HI, 2021.