73rd International Astronautical Congress (IAC), Paris, France, 18-22 September 2022.

IAC-22,A6,IP,x70583

Efficient Sensor Tasking for Space Situational/Domain Awareness

Cameron DeJac[a], Trevor Wolf[b], Josh Neel[a], Brandon Jones[b], Neil Dhingra[a]*, Alex Herz[a], Ella Herz[a]

[a] *Orbit Logic, Greenbelt, Maryland, USA 20770*
[b] *University of Texas, Austin*
* Corresponding Author neil.dhingra@orbitlogic.com

Abstract

The proliferation of space objects and the increasingly contested nature of space makes the Space Situational/Domain Awareness (SSA/SDA) problem important to enable effective space operations and to prevent collisions between space objects. With the increased burden placed on sensor networks performing SSA/SDA missions, it is important that sensing assets are used efficiently to track high-priority objects and maintain a high quality catalogue. Orbit Logic's Heimdall SSA/SDA tasking software performs intelligent sensor network tasking so that operators can efficiently obtain high-quality tracks on objects they care about while maintaining a high-quality catalogue. Heimdall creates cooperative schedules for ground- and space-based sensors to optimize an SSA/SDA-specific figure of merit (FOM), that reflects catalogue and/or mission objectives, while obeying user-specified planning 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. The FOM components and constraints can relate to sensors and/or collection data gathered on space object(s), identified by NORAD ID or a provided ephemeris, for factors such as observation timing (within specific time windows, around orbital events, or at a given cadence), sensor phenomenologies, sensor preferences, viewing geometry, weather, viewing conditions (e.g., object apparent visual magnitude), clearing radius, collection duration, radar cross section, other sensor parameters, and, now, expected orbit accuracy at given time(s) given the current track and planned observations. To ensure coverage and prevent redundancy, Heimdall can ingest plans from external sensor networks so that the Heimdall operator can condition their sensor scheduling around the data they expect to receive from partners. In addition, Orbit Logic's Heimdall SDA sensor tasking software now supports planning for expected orbit accuracy after data fusion is performed, allowing operators to specify required levels of expected orbit accuracy for space object tracks at user-specified times. Orbit accuracy is quantified by a user-specified function (e.g., the determinant) of the expected posterior track error covariance matrix obtained when the current track and data to be obtained from planned sensor collects are fused using a specified filter (e.g., the Extended Kalman Filter) by a module developed by the University of Texas, Austin. This can be leveraged to achieve high-accuracy tracks on high-value targets more efficiently. In addition, it can be used to intelligently reduce sensing on low-priority objects so that sufficient accuracy is maintained with fewer collects.

Keywords: Space Situational Awareness, Space Domain Awareness, Sensor Scheduling, Sensor Resource Management, Value of Information, Mission Planning and Scheduling

Acronyms/Abbreviations

Application Programming Interface (API)
Battle Management, Command, and Control (BMC2)
Collection Planning and Analysis Workstation (CPAW)
Combined Space Operations Center (CSpOC)
Department of Defense (DoD)
Electro-Optical (EO)
Extended Kalman Filter (EKF)
Figure of Merit (FOM)
Graphical User Interface (GUI)
Infrared (IR)
Integrated Sensor Support Plan (ISSP)
North American Aerospace Defense Command (NORAD)
Resident Space Object (RSO)
Space Domain Awareness (SDA)
Space Situational Awareness (SSA)
Space Surveillance Network (SSN)
Transmission Control Protocol/Internet Protocol (TCP/IP)
Value of Information (VOI)

1. Introduction

The proliferation of space objects and their increasingly complicated behaviours means that high-quality, often time-sensitive, tracks must be maintained on many objects for effective Space Situational/Domain Awareness (SSA/SDA). Such elevated requirements impose a larger burden on sensor networks. There has been much recent activity centred around fostering collaboration between the government and non-traditional sensor networks operated by commercial or international partners. Such partnerships seek to allow exquisite government sensors to focus on high-priority or sensitive tasks to ensure coverage – so the entire space catalogue is sufficiently monitored and so all tasking requirements are satisfied – and to limit redundancy – so that time from high-value sensing assets is not used when time from cheaper assets could be used to complete the same mission.

However, more sensing hardware and data sharing agreements themselves are not sufficient; intelligent sensor tasking is required to unlock the potential that these partnerships hold. To reduce the burden on exquisite government sensors, their planned tasking must be conditioned on partner sensor plans. To improve key catalogue health metrics, plans must be generated that target these metrics directly.

Orbit Logic's Heimdall SDA tasking software performs intelligent sensor network tasking that empowers such government-commercial and international SDA partnerships to impact space catalogue health. Heimdall cooperatively tasks ground- and space-based sensors to optimize an SDA-specific figure of merit (FOM) that reflects catalogue and/or mission objectives. Heimdall runs several optimization algorithms in parallel to generate different sensor schedules, and plans the schedule to optimize the FOM. Recently, Orbit Logic demonstrated Heimdall's upgraded capability to ingest commercial provider plans from the LeoLabs, Inc. and Numerica Corporation (now Slingshot Aerospace) commercial sensor networks and to output them to the Integrated Sensor Support Plan (ISSP) format – allowing the government to condition their sensor scheduling on partner plans.

Moreover, Heimdall supports orbit accuracy-based intelligent sensor tasking for SDA to directly plan for high quality tracks, rather than to plan for plan factors such as high revisit rates which are presumably correlated with high quality tracks. 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 Value of Information (VOI) and omitting tasks that are expected to provide low 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 manoeuvre 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 fulfil. 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.


Figure 1: Heimdall Logo

Alternatively, instead of in strict constraints, these factors can be used in components in the configurable SDA-specific figure of merit (FOM) to reflect catalogue 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.

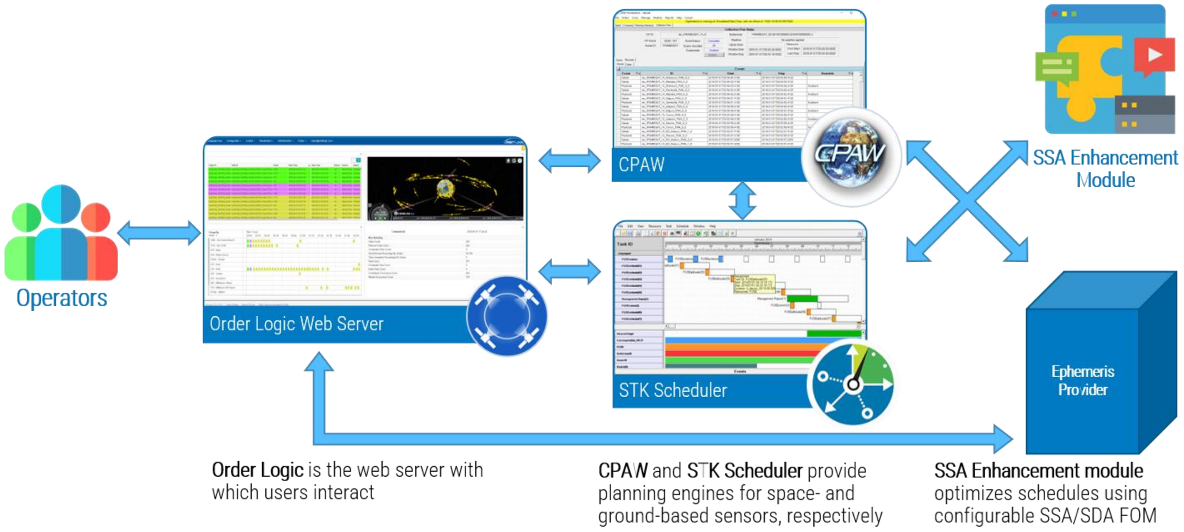

Figure 2: Heimdall System Architecture Diagram: Heimdall builds on and enhances mature Orbit Logic products to create optimized SSA/SDA sensor schedules

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 expected a posteriori estimation error covariance matrix and can correspond to physical properties of the covariance ellipse, such as the volume or major axis length.

1.1 Background

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. Heimdall can navigate this trade off via differing order specifications on each object.

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 manoeuvres, 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.

2. Heimdall for Optimized Sensor Scheduling in Support of Complex SSA/SDA Missions

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.

2.1 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 fulfilment status through the shared Object Catalogue 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 Application Programming Interface (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 Catalogue 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. Fulfilment status based on planned observations is stored in the Heimdall Object Catalogue 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.

2.2 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 behaviour 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.

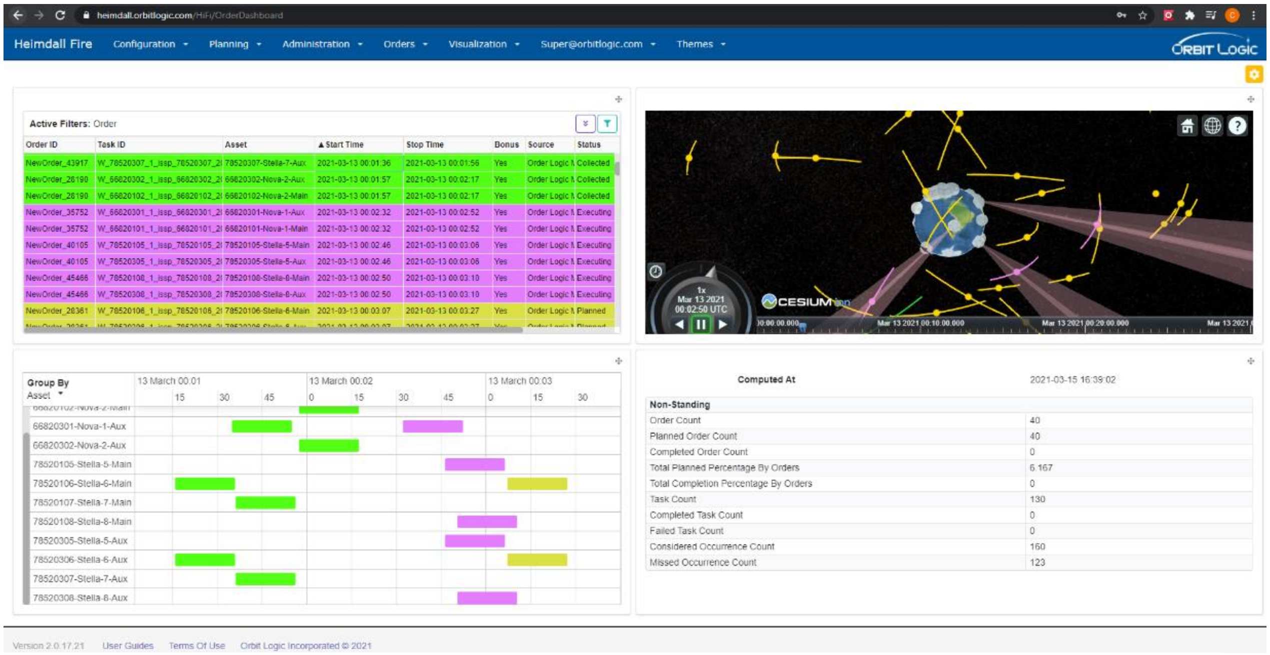

Active Filters: Order

Order ID	Task ID	Asset	▲Start Time	Stop Time	Bonus	Source	Status
NewOrder_43917	W_78520307_1_issp_78520307_2	78520307-Stella-7-Aux	2021-03-13 00:01:36	2021-03-13 00:01:56	Yes	Order Logic N	Collected
NewOrder_28190	W_66820302_1_issp_66820302_2	66820302-Nova-2-Aux	2021-03-13 00:01:57	2021-03-13 00:02:17	Yes	Order Logic N	Collected
NewOrder_28190	W_66820102_1_issp_66820102_2	66820102-Nova-2-Main	2021-03-13 00:01:57	2021-03-13 00:02:17	Yes	Order Logic N	Collected
NewOrder_35752	W_66820301_1_issp_66820301_2	66820301-Nova-1-Aux	2021-03-13 00:02:32	2021-03-13 00:02:52	Yes	Order Logic N	Executing
NewOrder_35752	W_66820101_1_issp_66820101_2	66820101-Nova-1-Main	2021-03-13 00:02:32	2021-03-13 00:02:52	Yes	Order Logic N	Executing
NewOrder_40105	W_78520105_1_issp_78520105_2	78520105-Stella-5-Main	2021-03-13 00:02:46	2021-03-13 00:03:06	Yes	Order Logic N	Executing
NewOrder_40105	W_78520305_1_issp_78520305_2	78520305-Stella-5-Aux	2021-03-13 00:02:46	2021-03-13 00:03:06	Yes	Order Logic N	Executing
NewOrder_45466	W_78520108_1_issp_78520108_2	78520108-Stella-8-Main	2021-03-13 00:02:50	2021-03-13 00:03:10	Yes	Order Logic N	Executing
NewOrder_45466	W_78520308_1_issp_78520308_2	78520308-Stella-8-Aux	2021-03-13 00:02:50	2021-03-13 00:03:10	Yes	Order Logic N	Executing
NewOrder_28361	W_78520106_1_issp_78520106_2	78520106-Stella-6-Main	2021-03-13 00:03:07	2021-03-13 00:03:27	Yes	Order Logic N	Planned

Computed At 2021-03-15 16:39:02

Non-Standing	
Order Count	40
Planned Order Count	40
Completed Order Count	0
Total Planned Percentage By Orders	6.167
Total Completion Percentage By Orders	0
Task Count	130
Completed Task Count	0
Failed Task Count	0
Considered Occurrence Count	160
Missed Occurrence Count	123

Figure 3: Heimdall Dashboard Page with the light color theme

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.

2.3 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 catalogue 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 modelling. 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.

2.4 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.

2.5 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 catalogue 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.

2.6 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 API configuration commands to some of Heimdall solution component applications.

2.7 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 catalogues.

3. Efficient Planning and Scheduling with Heimdall

To demonstrate the efficacy of Heimdall, we present three studies. In the first, we consider the impact of track accuracy-based tasking on sensor network efficiency. In the second, we demonstrate how Heimdall's coordination enhances cooperation between sensor networks. Finally, we show how Heimdall can rapidly adjust tasking in response to encountered events.

3.1 Track accuracy-Based Tasking

We consider tasking on a particular real-world space object with three different sensors. 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.

3.1.1 Better results for similar investment

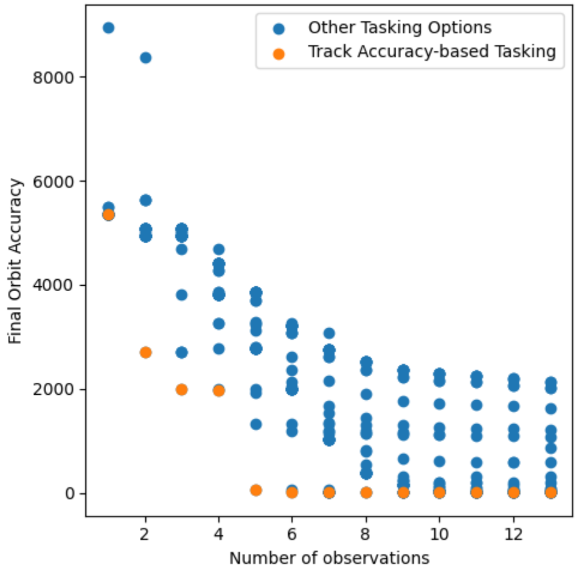

Figure 4:Track accuracy (trace of the a posteriori estimation error covariance matrix) at a target time for different potential sensor observation schedules for a specific RSO.

To compare the trade off 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 4 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.

3.1.2 Lower investment for similar results

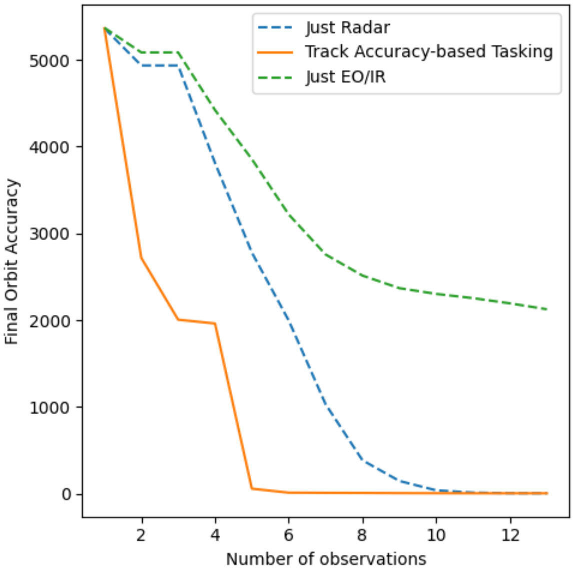

Figure 5: Track accuracy at a target time as more observations are fused. The observations' timing is identical.

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 5 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 trade off 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 catalogue maintenance mission.

3.2 Coordination of Cooperative Sensor Networks

We consider the importance of coordinating multiple sensor networks for SDA in the context of providing persistent coverage. Because large gaps between successive observations leaves room for an object to manoeuvre unexpectedly or to drift out of its nominal orbit, we specify a required revisit rate for maintaining sufficient awareness of an object. Note that the phase of the observations is left up to the algorithms (e.g., if an object must be observed every hour, it can be observed every x:15 or x:20 or other time) and observations can be *more* frequent than desired (e.g., if it must be observed every hour, the object is observed at 12:15, 1:00, 2:00, 2:10) and still fulfil the required revisit rate.

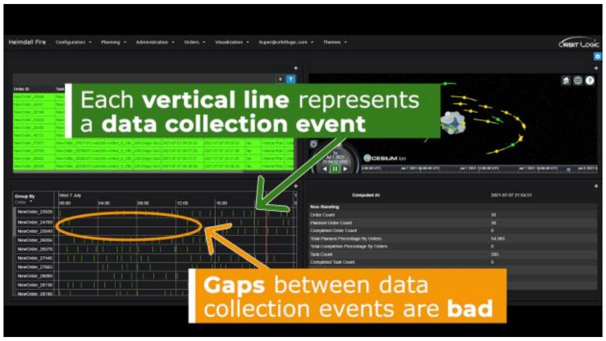

Figure 6: Heimdall Screenshot indicating gaps in coverage.

The performance metric for this study is the percent of desired regularly spaced observations that are fulfilled. Desired regularly spaced observations occur when a successive observation is scheduled before the desired revisit period has expired and multiple desired regularly spaced observations are deemed unfulfilled when multiple successive revisit periods have elapsed, i.e., if the desired revisit rate is 1 hour and gap between observations is 1:10, one desired observation is deemed missed and if it is 2:01, 2 desired observations are deemed missed.

We consider the case where two sensor networks are attempting to observe the same set of objects with the same constraints. In this scenario, both sensor networks are alone incapable of satisfying the requirement. We study the performance metric when each network tasks on the objects individually and delivers data separately to a central data repository in comparison with when the networks coordinate tasking.

In this scenario, two sensor networks are delivering observations on a common set of objects to a common repository where data will be fused. For awareness of manoeuvres, each object must be observed every hour. Note that this is a requirement distinct from a requirement related to orbit accuracy, since it is only concerned with observation frequency and not value of information (e.g., from geometry, sensor type, etc.).

One sensor network has the capacity to perform 50.926% of the desired observations while the other has the capacity to perform 59.722% of the desired observations. When their plans are uncoordinated and the union of their observations is taken, there is overlap in delivering desired regularly spaced observations and many gaps persist; together and uncoordinated, the sensor networks deliver 77.778% of the desired observations. When Heimdall coordinates the sensing, it is aware of overlapping coverage and gaps in coverage and so the joint plan for both sensor networks delivers 98.148% of the desired observations.

Clearly, Heimdall coordinating sensor schedules delivers a large increase in performance. From the perspective of the second network, buying data from the first network would take coverage from 59.722% to 98.148% instead of to 77.778%. This means that the money spent on buying data from the first network would go 2.128 times as far! Overall, Heimdall's coordination results in 26.19% more coverage/tasking than when the sensor networks are uncoordinated.

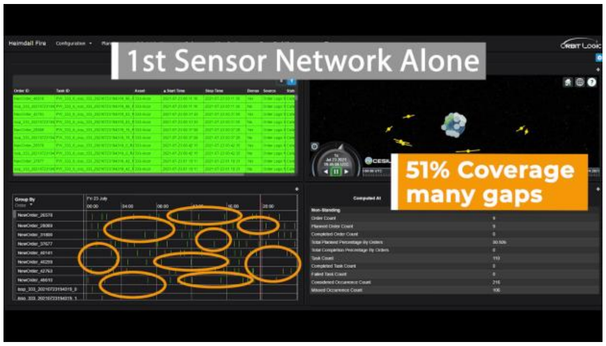

Figure 7: Heimdall screenshot showing the performance of Sensor Network 1, which completes 50.926% of the desired tasks.

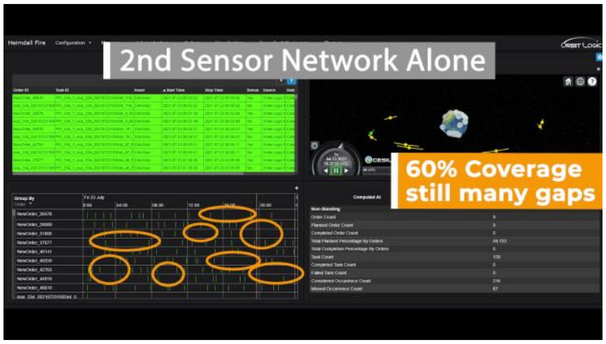

Figure 8: Heimdall screenshot showing the performance of Sensor Network 2, which completes 59.722% of the desired tasks.

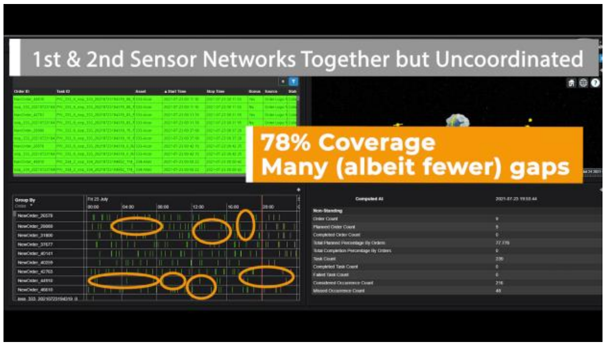

Figure 9: Heimdall screenshot showing the performance of Sensor Networks 1 and 2 together but uncoordinated, which completes 77.778% of the desired tasks.

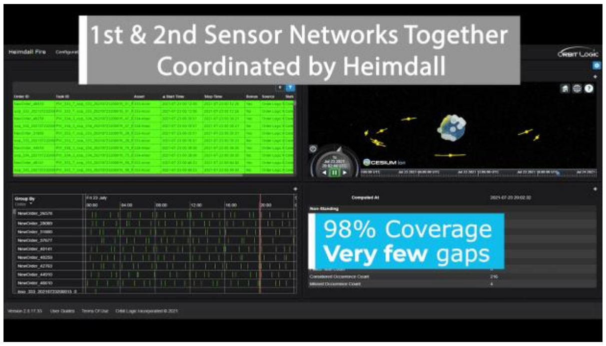

Figure 10: Heimdall screenshot showing the performance of Sensor Networks 1 and 2 when Heimdall coordinates the planning, which completes 98.148% of the desired tasks.

3.3 Heimdall for Reactive Retasking

We also consider a dynamic scenario in which a few objects are deemed worthy of elevated tasking frequency and tasking should be updated. This demonstrates the importance of fast algorithms and replanning.

In this scenario, the nominal plan is generated by Heimdall and coordinates the two sensor networks in the prequel. However, analysts or automated software deem the top two objects worth observing twice as often due to an elevated risk of manoeuvre or another factor. The bottom four objects are deemed a lower priority than the other objects, perhaps because they are at a *low* risk to manoeuvre. In summary: the top two objects are high priority and should be monitored every 30 minutes, the bottom four objects are low priority but should be monitored every hour, and the remaining objects are medium priority and should be monitored every hour. In the earlier scenario, this was not an issue because there was sufficient capacity to support full tasking across objects of all priority levels.

Heimdall adjusts planning accordingly, shifting observations from the low priority objects to support elevated tasking on the high priority objects. Without re-planning, only half of the desired high priority observations are fulfilled while all of the low-priority observations are fulfilled. Heimdall's dynamic replanning shifts this so that all of the high-priority observations are fulfilled and half of the low-priority observations are fulfilled.

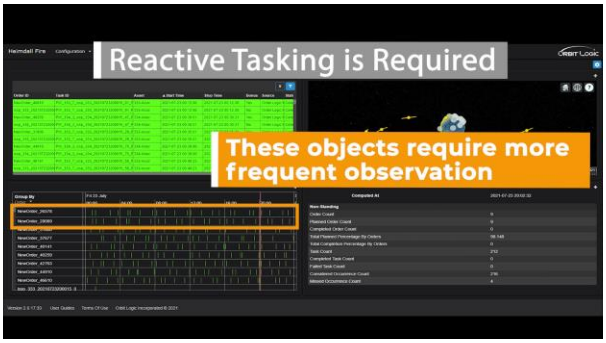

Figure 11: Heimdall screenshot showing the baseline performance where additional tasking is desired on the top two objects.

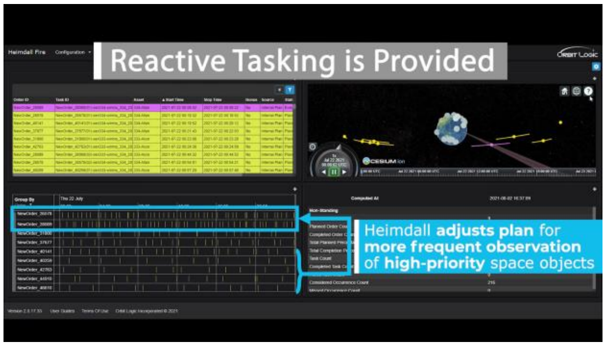

Figure 12: Heimdall screenshot showing the adjusted performance where additional tasking is performed on the top two objects by shifting tasking away from the lower four objects, which had been deemed lower priority.

4. Conclusions

The intelligent planning and scheduling capability provided by Heimdall optimizes information gain; coordinates an entire heterogeneous network across ground and space and with multiple phenomenologies; balances search, track, and other tasks; and enables rapid reaction via responsive retasking. We demonstrated how Heimdall increases sensor network efficiency when targeting track accuracy metrics, when coordinating multiple sensors or sensor networks, and when reacting to events that change planning parameters.

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.