

Maximizing the Utility of Non-Traditional Sensor Network Data for SDA

Neil K Dhingra, Cameron DeJac, 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

We present selected studies demonstrating how intelligent tasking maximizes the utility of non-traditional provider data for Space Domain Awareness (SDA). We compare gathered SDA sensor data under different tasking strategies: when sensor networks' tasking is/is not coordinated with one another (coordination more than doubles the impact of commercial SDA data purchases), when it does/does not target specific objectives (intelligent tasking reduces the size and rate of growth of object orbit estimation error covariance), and the impact of different algorithms on tasking performance (over seven consecutive planning periods, four different algorithms returned the best plan). These studies were enabled by Orbit Logic's Heimdall SDA sensor tasking software, which has recently been upgraded to ingest commercial provider data and for catalog-health-based tasking, features which can reduce the burden on exquisite government assets to ensure coverage and limit redundancy.

1. INTRODUCTION

The growing proliferation of space objects, and the increasingly contested nature of space as a warfighting domain, make the Space Domain Awareness (SDA) mission more difficult because more objects must be monitored and because effectively monitoring the emerging diverse types of space objects requires more sensor tasking with more complicated requirements. To alleviate the larger burden that these considerations impose on government sensors, there has been much recent activity centered 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 catalog 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 catalog 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 catalog health. Heimdall cooperatively tasks ground- and space-based sensors to optimize an SDA-specific figure of merit (FOM) that reflects catalog 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 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.

Orbit Logic has also partnered with the University of Texas at Austin to enhance Heimdall for sensor scheduling that targets specific catalog improvement metrics, including those related to the predicted error uncertainty associated with different objects. This would enable operators to task sensors around the estimated accuracy of object orbit estimates, e.g., to ensure that the maximum error covariance across a catalog is below a certain threshold. Moreover, these requirements can be heterogeneous, e.g., so that objects in crowded orbits have better estimates, and they can interact with other requirements, e.g., so that objects that frequently maneuver are still persistently monitored at a minimum observation frequency.

To quantify the effect of tasking methodologies on space catalogs, we simulated data collection by model sensor networks representative of government and partner networks in different scenarios with different sensor tasking strategies. These include static scenarios, in which space objects behave nominally, as well as dynamic scenarios, in which observed actions, e.g., satellite maneuvers, prompt revised tasking requirements and updated scheduling to satisfy those requirements, e.g., elevated collection frequency on maneuvering objects. We simulated data collection using the different sensor schedules with a non-zero probability of missed detection. To generate space catalogs from these different sets of simulated data, we updated the object orbit estimates assuming perfect data association; the effect of inaccurate data association may be studied in follow on work. By deriving and comparing metrics from these resulting space catalogs, we illustrate the effect of different tasking strategies on catalog health.

The different sensor tasking strategies are: when government sensor network tasking is and is not conditioned on partner sensor network tasking, when tasking does and does not target specific mission objectives (e.g., limiting the catalog-wide maximum error uncertainty below a threshold), and when different optimization algorithms are used for generating task plans. If government tasking is not conditioned on partner plans, there is a risk that partner-provided data is redundant with government-collected data and that not all tasks will be fulfilled. If tasking does not target specific mission objectives, there is a risk that sensor time is poorly used on collection tasks that yield low value-of-information data. Finally, the sensor scheduling problem itself is very difficult – even simplified versions of it are NP hard – so it is unlikely that any one sensor tasking algorithm will always generate effective plans. We demonstrate how different algorithms excel for different mission scenarios and how running several algorithms in parallel and choosing the best generated sensor schedule (as measured by an SDA-specific FOM) enhances the quality of the collected data.

The studies we present are enabled by Orbit Logic's Heimdall Sensor Tasking software for SDA. Heimdall performs intelligent sensor scheduling and facilitates effective partnerships between the government and non-traditional data providers.

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.

Of course, requirements on data quality and quantity are coupled with one another, with the object, and with the sensors collecting data. Objects more prone to maneuvers may require more frequent data collection. 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.

In the interest of simplicity, we isolate particular tasking challenges and study them separately to illustrate and quantify the importance of intelligent planning and scheduling. In particular, we study the challenges of 1) achieving sufficiently numerous and regular observations across multiple sensor networks 2) achieving high quality tracks on an object and 3) the fundamental NP hard complexity of tasking problems. We discuss these further in Section 4. 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.

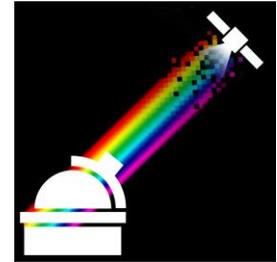


Figure 1: Heimdall Logo

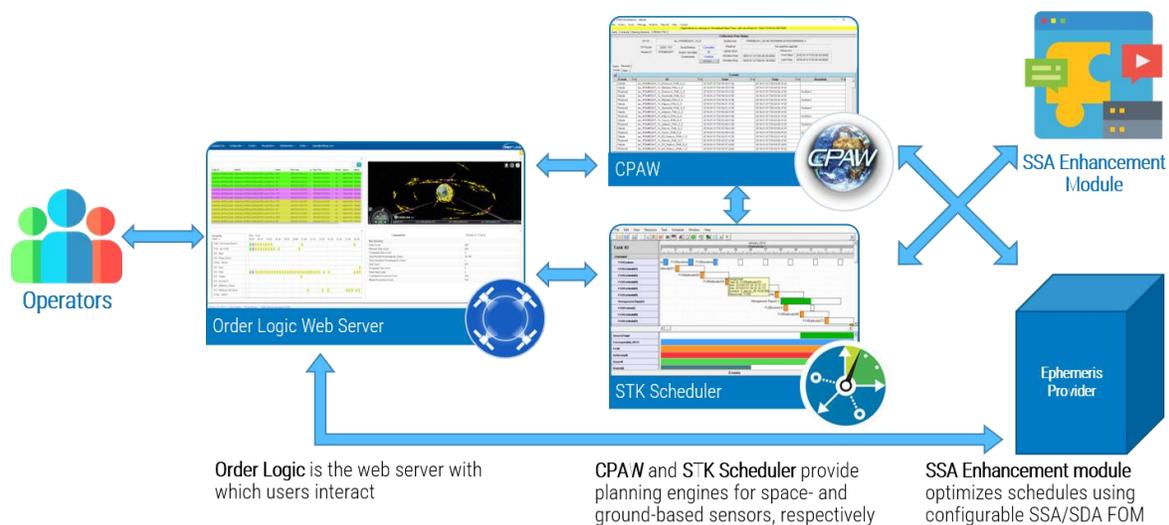


Figure 2: Heimdall System Architecture Diagram: Heimdall builds on and enhances mature Orbit Logic products to create optimized SSA/SDA sensor schedules

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.

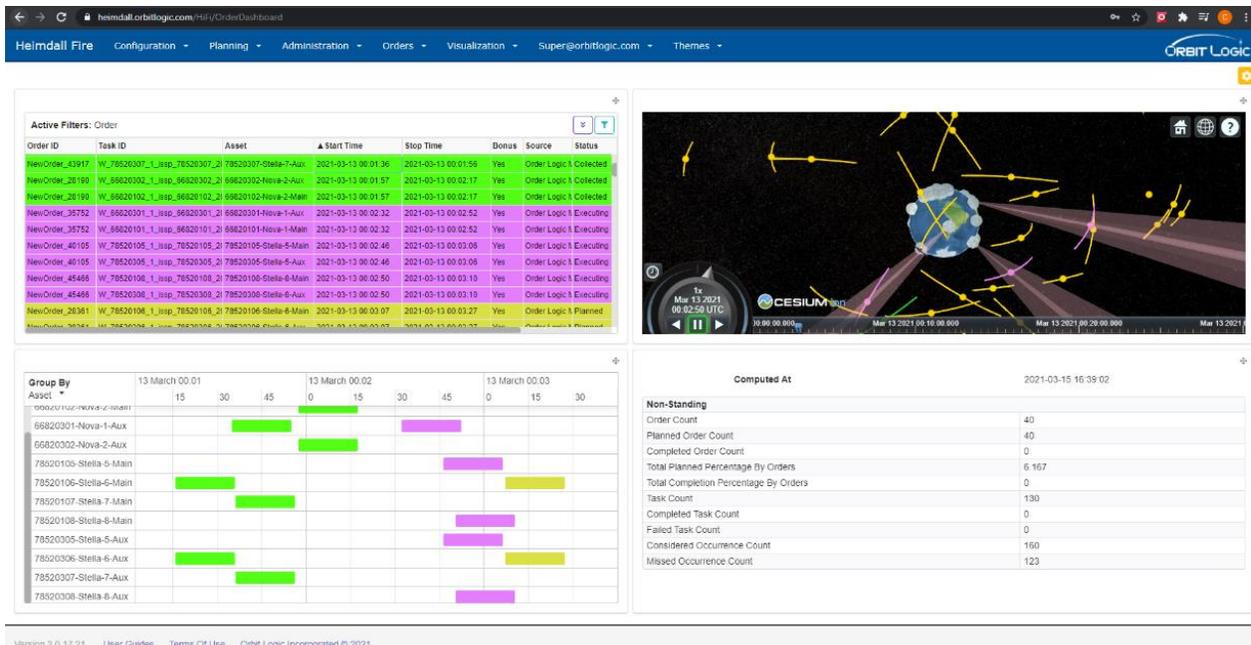


Figure 3: Heimdall Dashboard Page with the light color theme

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). 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-site 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. DESCRIPTION OF STUDIES AND PERFORMANCE METRICS

In this paper, we study the effect of three specific tasking challenges on the quality of the catalog: 1) achieving sufficiently numerous and regular observations across multiple sensor networks 2) achieving high quality tracks on an object and 3) the fundamental NP hard complexity of tasking problems. We study these challenges separately and with limited, small-size problems to provide an easily-understandable basis for understanding the complex mechanisms in play when performing planning and scheduling for SDA. These simplifications are made for the ease of exposition and comprehension. As tasking challenges interact in more complex ways, the need for better planning and scheduling is more pronounced.

Future work will focus on more comprehensive studies with more complexity including larger catalogs, more sensor networks, more sensors, more tasking requirements, and so on. As discussed in Section 5, we see stark differences in performance with different tasking *even in these limited studies*, highlighting the importance of intelligence planning and scheduling.

Study 1: Persistent Observations Across Different Sensor Networks

The first study quantifies 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 maneuver 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 fulfill the required revisit rate.

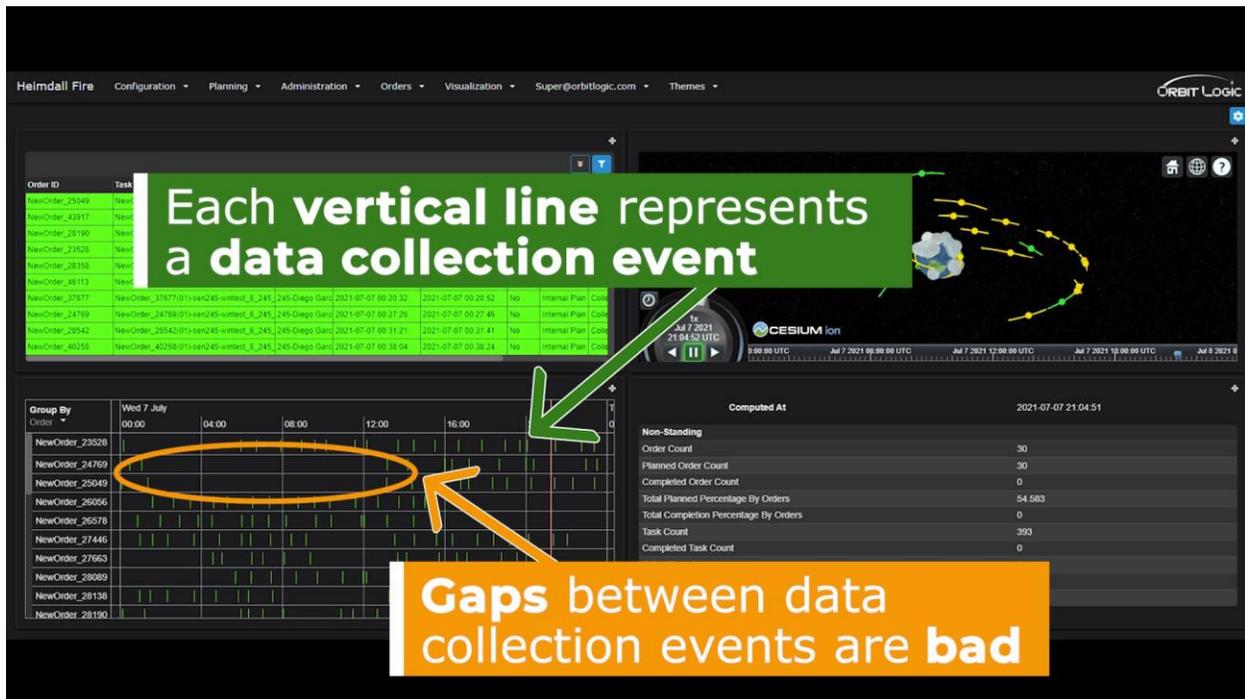


Figure 4: Heimdall Screenshot indicating gaps in coverage.

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.

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

Study 2: Tasking to Target Track-Specific Metrics

Although regularly spaced observations are important to detect object maneuvers, during nominal operations it is more effective to task to optimize track-specific metrics. These metrics are typically based on the expected estimation error covariance of the tracked object conditioned on an estimation plan. Clearly, this will depend on the filter used to fuse information for different sensors. For this demonstration, we consider an Extended Kalman Filter and a plan with sensors akin to a radar and to an optical sensor.

Study 3: Sensitivity of Algorithm Performance to Scenario Parameters

Planning and scheduling problems are typically NP hard and this is true in general for SDA sensor tasking. As both the Travelling Salesperson and Bin Packing Problems are NP hard, SDA sensor scheduling – which includes these problems as special instances – also has this property.

As such, we cannot expect any practical algorithm to reliably do well. Any polynomial time algorithm will be a heuristic and may perform poorly on some cases. We note that there are certain special cases of Travelling Salesperson and Bin Packing Problems that admit heuristics with guaranteed performance levels – such as that based on a Minimum Spanning Tree for the Travelling Salesperson Problem when the distances obey the triangle inequality – but the required assumptions are not true in general for SDA. For example: because observation targets are objects in space that are moving, that the triangle inequality is not observed; there are observation windows that preclude task fulfillment when the object is not in view or when the sun interferes with sensors, etc.; there are multiple sensors that must be coordinated so there are elements of a Vehicle Routing Problem; not every task may be fulfilled so there are elements of an Orienteering Problem. The simplest problem formulation that well describes sensor scheduling is a Time Dependent Team Orienteering Problem with Time Windows (TDTOPTW), but even this complicated formalism does not capture the full complexity of SDA tasking requirements.

It is therefore to be expected that different algorithms perform best for different planning scenarios. The performance metric we use to compare sensor plans is the SDA-specific FOM used in Heimdall as the objective function. The parameters of the FOM are not public, but we compare algorithm performance in terms of the percent performance relative to the highest scoring algorithm.

We study different scenarios with the *same* set of objects, the *same* requirements, and the *same* set of sensors, and vary *only* the particular day of the observations. Even in this scenario where most factors are held constant, different algorithms provide better performance on different days. This motivates Heimdall's operational standard to run these algorithms in parallel and choose the best scoring plan.

Since this study is concerned with algorithm performance at the time of planning, we do not do separate studies for static and dynamic scenarios here. From the perspective of the planning algorithms, dynamic scenarios, where planning must be re-run in response to events, are the same as different static scenarios; they differ only in problem data for the planning and scheduling algorithms.

5. RESULTS OF STUDIES

We now explain the results of our studies. In summary: coordinating sensor networks with Heimdall more than doubles the impact of money spent on commercial SDA data and allows for efficient dynamic replanning; intelligent tasking reduces the size and rate of growth of object orbit estimation error covariance; and four different algorithms created the best schedules on seven different scenarios, motivating a concept of operations in which multiple algorithms are run in parallel to create several candidate schedules from which the best is chosen.

Study 1: Persistent Observations Across Different Sensor Networks *Static Scenario*

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 maneuvers, 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 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.

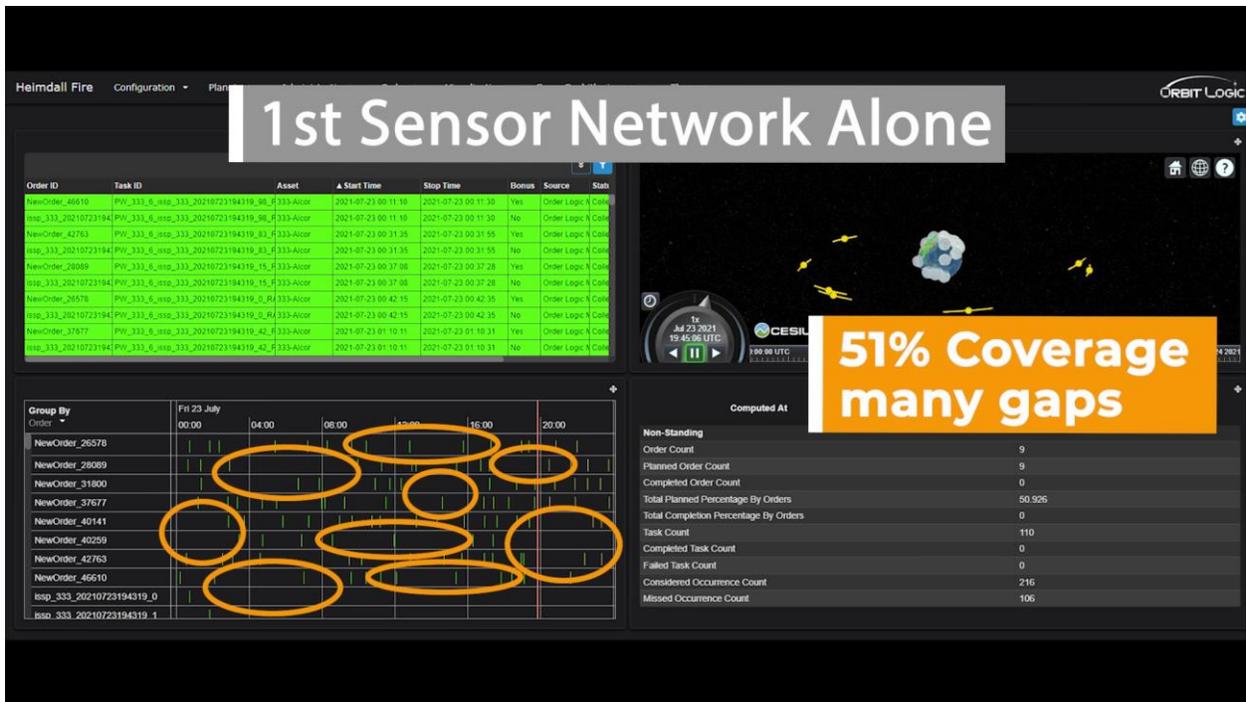


Figure 5: Heimdall screenshot showing the performance of Sensor Network 1, which completes 50.926% of the desired tasks.

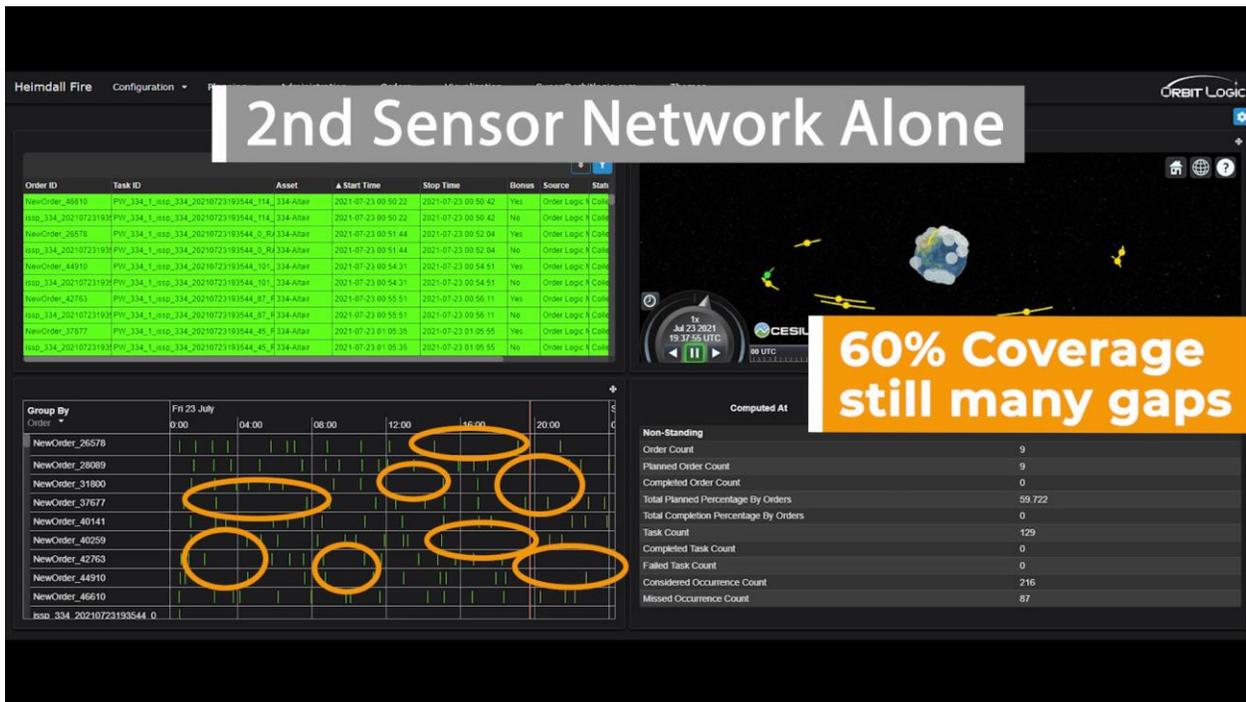


Figure 6: Heimdall screenshot showing the performance of Sensor Network 2, which completes 59.722% of the desired tasks.

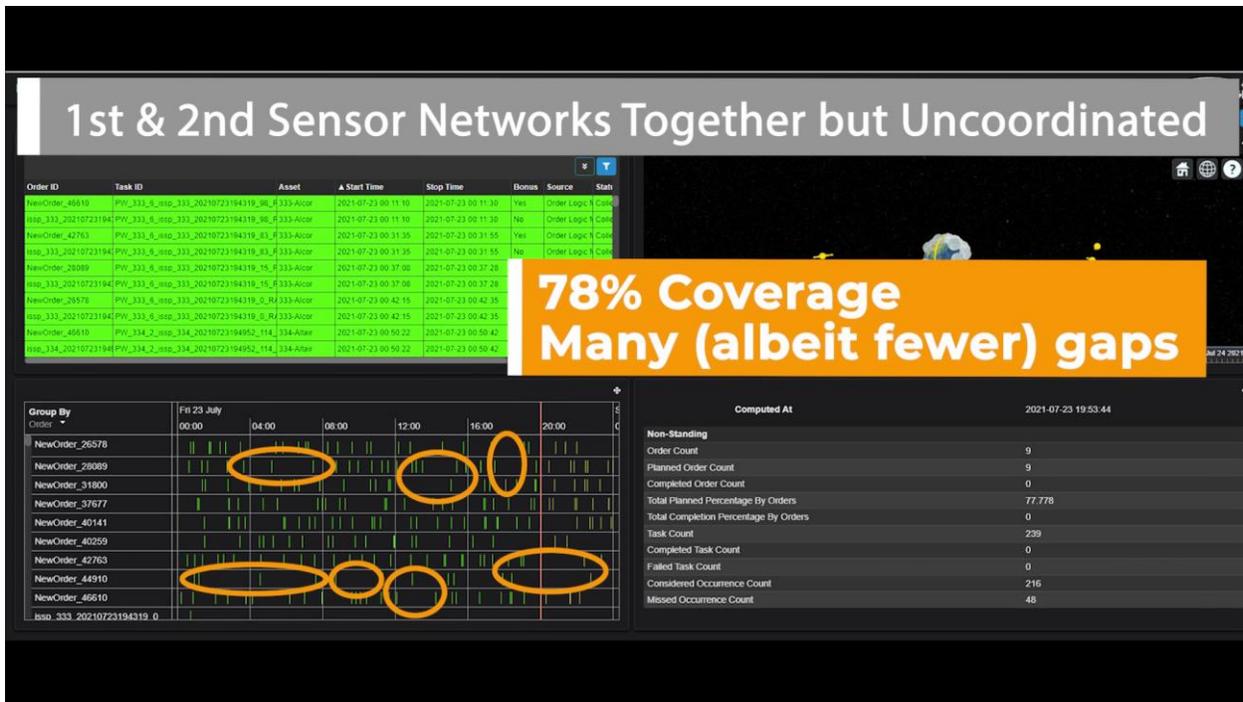


Figure 7: Heimdall screenshot showing the performance of Sensor Networks 1 and 2 together but uncoordinated, which completes 77.778% of the desired tasks.

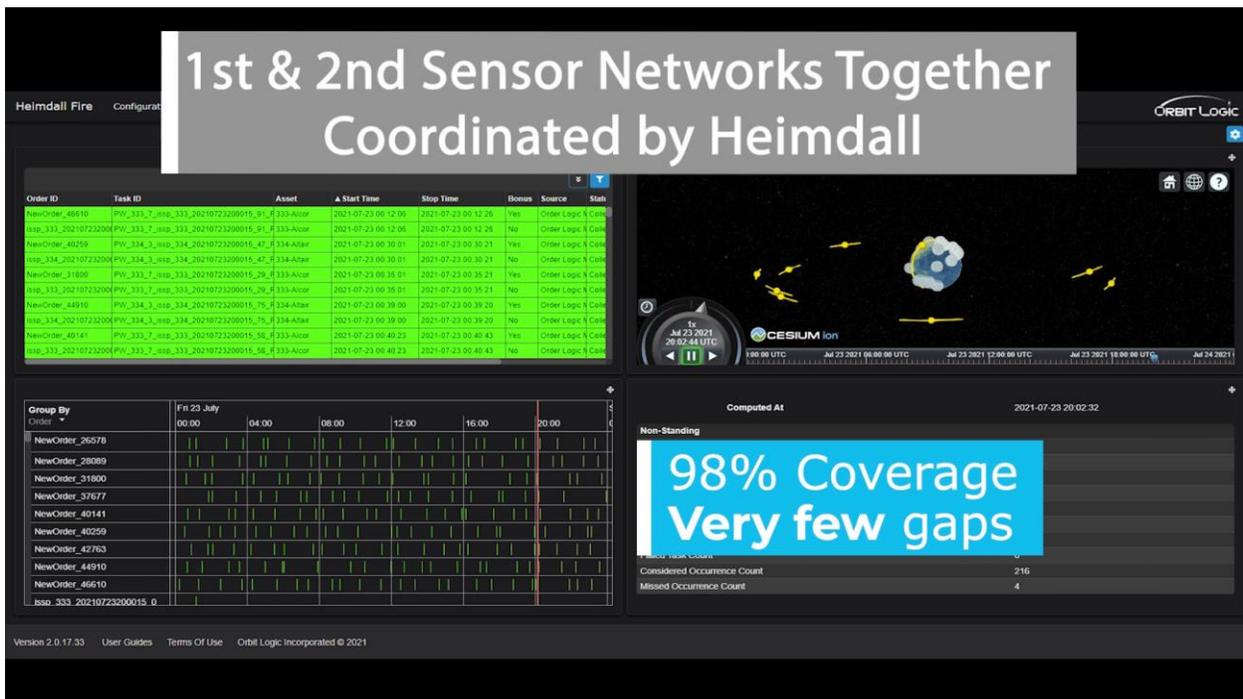


Figure 8: Heimdall screenshot showing the performance of Sensor Networks 1 and 2 when Heimdall coordinates the planning, which completes 98.148% of the desired tasks.

Dynamic Scenario

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 maneuver 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 maneuver. 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.

Objects	Coverage with Original Plan	Coverage with Updated Plan
High Priority	50%	100%
Medium Priority	100%	100%
Low Priority	100%	50%

Table 1: Coverage With and Without Replanning.

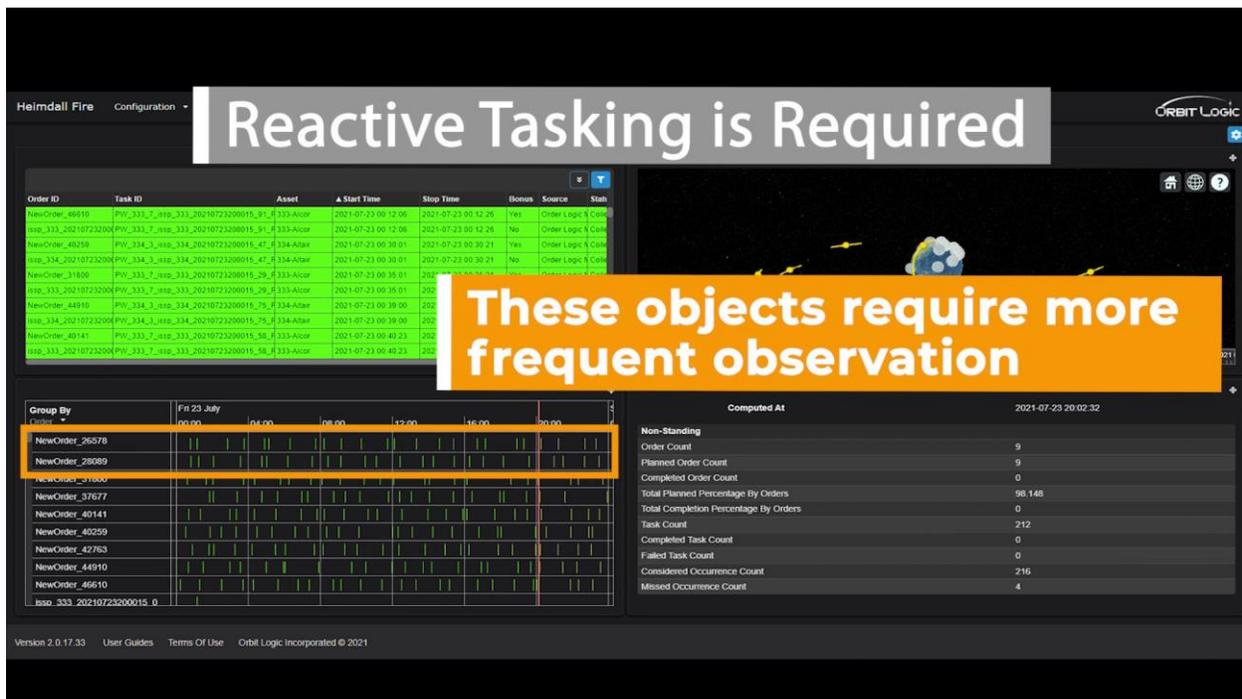


Figure 9: Heimdall screenshot showing the baseline performance where additional tasking is desired on the top two objects.

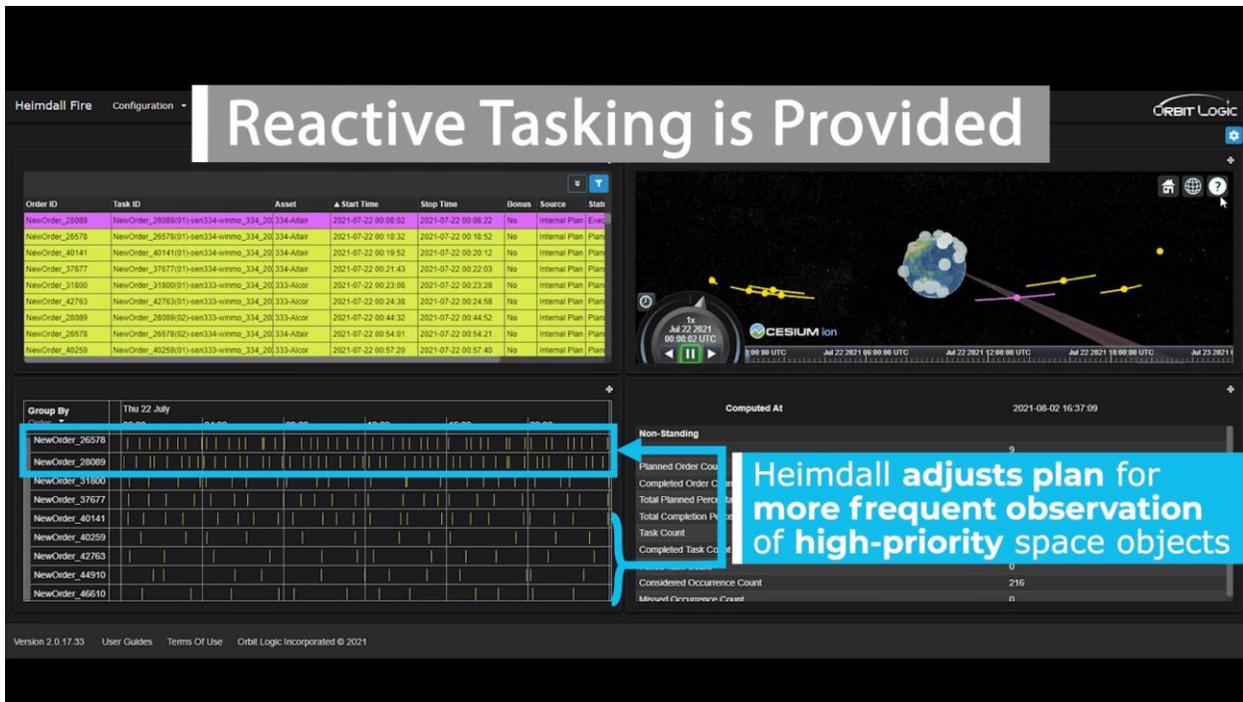


Figure 10: Heimdall screenshot showing the adjusted performance where additional tasking is performed on the top two objects by shifting tasking away from the lower four objects, which had been deemed lower priority.

Study 2: Tasking to Target Track-Specific Metrics

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 are less regular but occur with a mix of the two sensors.

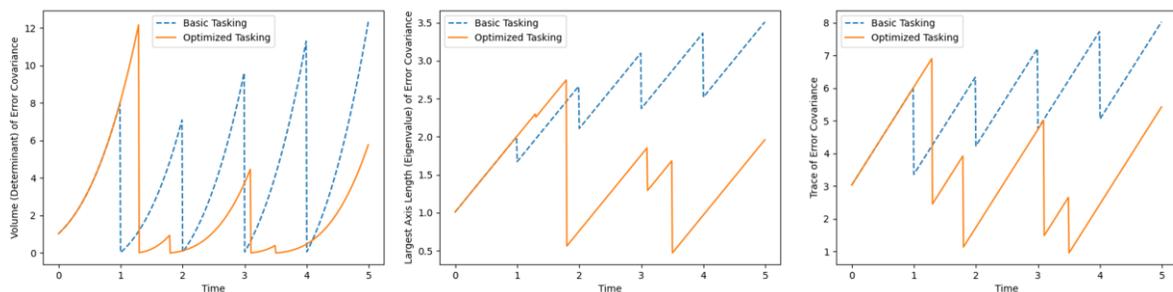


Figure 11: Effect of track metric-optimized sensor tasking.

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 volume (i.e., determinant), 1.79 by semimajor axis (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.

Study 3: Sensitivity of Algorithm Performance to Scenario Parameters

In this study, we examined how different algorithms performed on different problems. To emphasize the sensitivity of algorithms to the problem data, we used the same algorithms on the same set of objects with the same order requirements and the same sensors for a twenty-four hour period, only varying which period was considered. All of the sensors considered are space-based and this study deals exclusively with space-to-space observations. It is worth noting that although some of these algorithms use a random seed, many are deterministic, meaning that the only difference in performance arises from the change of day, e.g., due to small changes in orbit timing or other similar factors.

We quantify algorithm performance using the SDA-specific FOM – the quantity which they are trying to maximize – and discuss this quantity as the percent of the highest scoring FOM for that day; in Table 2 and Figure 12, the best scoring algorithm will always score a 1.

In seven different planning periods, four different algorithms delivered the best plan. Of the four algorithms that scored highest at least once, three of them also scored worst of those four on another day. No one algorithm was good enough to use exclusively. Since these algorithms are cheap to run relative to the cost of sensor time and analyst resources, it is worth running them all in parallel to ensure that sensor network resources can be used as optimally as possible. Finally, we note that planning problems will likely differ much more than they did in this study; additional space objects are continuously being deployed, more sophisticated capabilities are continuously being demonstrated on orbit, and more sensing capabilities are being continuously installed on ground stations across the Earth and on board satellites across different orbital regimes.

Normalized Algorithm Performance								
Algorithm	26-Aug	27-Aug	28-Aug	29-Aug	30-Aug	31-Aug	1-Sep	Worst
1	0.99786	0.999925	1	0.985045	0.991931	1	0.975359	0.975359
2	0.996181	1	0.996231	0.97627	0.983576	0.985187	0.968502	0.968502
3	0.997828	0.996673	0.991091	0.977028	0.978816	0.982122	0.968116	0.968116
4	0.99542	0.999727	0.996667	0.981891	0.980483	0.986257	0.962142	0.962142
5	1	0.999846	0.993644	0.976889	0.983044	0.979589	0.966913	0.966913
6	0.992774	0.997345	0.980362	1	1	0.995422	1	0.980362
7	0.994769	0.992354	0.983377	0.964077	0.979316	0.97788	0.958982	0.958982
8	0.984397	0.982694	0.985584	0.97349	0.986508	0.977662	0.962915	0.962915
9	0.96818	0.976166	0.977545	0.967233	0.968672	0.981012	0.951785	0.951785

Table 2: Performance of Planning Algorithms on Different Days with the Same Orders and Sensors. Each algorithm's FOM is shown as the percent of the maximum FOM across all algorithms for that day. Higher scores are better.

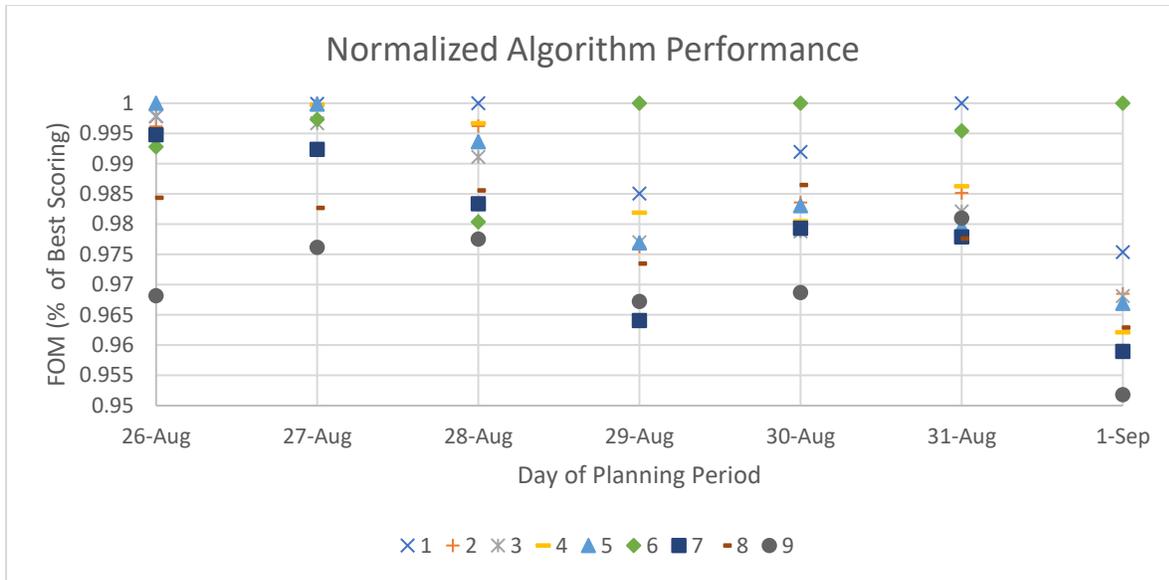


Figure 12: Performance of Planning Algorithms on Different Days with the Same Orders and Sensors. Each algorithm's FOM is shown as the percent of the maximum FOM across all algorithms for that day. Higher scores are better.

6. CONCLUSIONS

In this paper, we illustrate the importance of intelligent sensor tasking for SDA with simple, easy-to-understand examples that contain quite stark results. These studies use Orbit Logic's Heimdall SDA sensor tasking software to provide the intelligent tasking, and it makes a demonstrable impact on plan efficacy.

We first showed the importance of coordinating different sensor networks; using Heimdall more than doubles the impact of money spent on commercial SDA data and allows for efficient dynamic replanning. It is worth noting that Heimdall has recently been upgraded to ingest commercial provider data – from Numerica and LeoLabs – which is essential to do coordination effectively.

We then demonstrated that intelligent tasking to target track metrics based on the existing object orbit estimation error covariance, the sensor error covariances, and sensor-object geometries, makes a large impact on the final track metrics. It reduces the size and rate of growth of the estimation error covariance.

Finally, we studied the impact of different planning algorithms on the solution quality for this NP hard problem. On seven slightly different scenarios, four different algorithms scored highest. This result underscores the importance of running multiple planning algorithms in parallel to achieve good results. In operations, planning problems can easily evolve and be very different from the problems on which algorithms are initially evaluated, especially as the space domain evolves and becomes more complicated.

Although we have demonstrated the importance of intelligent tasking for effective SDA, we have presented studies in a way that is simple and easy to understand. This paper motivates further studies that specifically consider the impact of intelligent tasking on SDA with real sensor networks and real partnerships between US government, commercial, and international governments. Orbit Logic would like to pursue this avenue of research in conjunction with government and commercial partners.

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.