# Heimdall System for MSSS Sensor Tasking

**Alex Herz**
*Orbit Logic*
**Brandon Jones**
*University of Colorado*
**Ella Herz, Doug George**
*Orbit Logic*
**Penina Axelrad, Steve Gehly**
*University of Colorado*

## ABSTRACT

In Norse Mythology, Heimdall uses his foreknowledge and keen eyesight to keep watch for disaster from his home near the Rainbow Bridge. Orbit Logic and the Colorado Center for Astrodynamics Research (CCAR) at the University of Colorado (CU) have developed the Heimdall System to schedule observations of known and uncharacterized objects and search for new objects from the Maui Space Surveillance Site.

## 1. INTRODUCTION

Heimdall addresses the current need for automated and optimized SSA sensor tasking driven by factors associated with improved space object catalog maintenance. Orbit Logic and CU developed an initial baseline prototype SSA sensor tasking capability for select sensors at the Maui Space Surveillance Site (MSSS) using STK and STK Scheduler, and then added a new Track Prioritization Component for FISST-inspired computations for predicted Information Gain and Probability of Detection, and a new SSA-specific Figure-of-Merit (FOM) for optimized SSA sensor tasking. While the baseline prototype addresses automation and some of the multi-sensor tasking optimization, the SSA-improved prototype addresses all of the key elements required for improved tasking leading to enhanced object catalog maintenance.

The Heimdall software solution provides a configurable, automated system to improve sensor tasking efficiency and responsiveness for SSA applications. The FISST algorithms for Track Prioritization, SSA specific task and resource attributes, Scheduler algorithms, and configurable SSA-specific Figure-of-Merit together provide optimized and tunable scheduling for the Maui Space Surveillance Site and possibly other sites and organizations across the U.S. military and for allies around the world.

## 2. HEIMDALL SOLUTION

### Heimdall Mission Planning Operations

Heimdall scheduling provides a loop that includes modeling the system, creating priorities, creating a schedule, collecting observation data, and feeding the updated data back through the loop. Whether initiated manually, automatically on a set schedule, or based on trigger events, the proposed process will account for all critical factors every time, and create an optimized schedule based on a tunable Figure-of-Merit (FOM) that will help ensure the consistent achievement of system goals. These following goals have been identified for the Heimdall system:

- Maximize uncorrelated target (UCT) observation
- Minimize known object state covariance
- Track objects exhibiting uncharacteristic behavior
- Maintain custody of the known object catalog
- Maximize detection of unknown objects

Heimdall operations start with a catalog of objects including information such as last observation and special perturbation (SP) state vectors with covariance or two-line element (TLE) sets. Objects from this catalog along with sensor information are fed into the Modeling Environment for physical modeling including access and lighting computations. Objects from the Modeling Environment are pulled into the Heimdall Scheduling architecture for determining observation opportunities. The Heimdall Track Prioritization component uses a combination of observation data, object models, and observation opportunities to develop prioritization metrics for each target object and observation opportunity. The Heimdall Scheduling Algorithms then use the prioritization metrics and models along with a customized Figure-of-Merit to create an optimized, deconflicted observation schedule. Fig. 1, depicts the Heimdall system architecture.

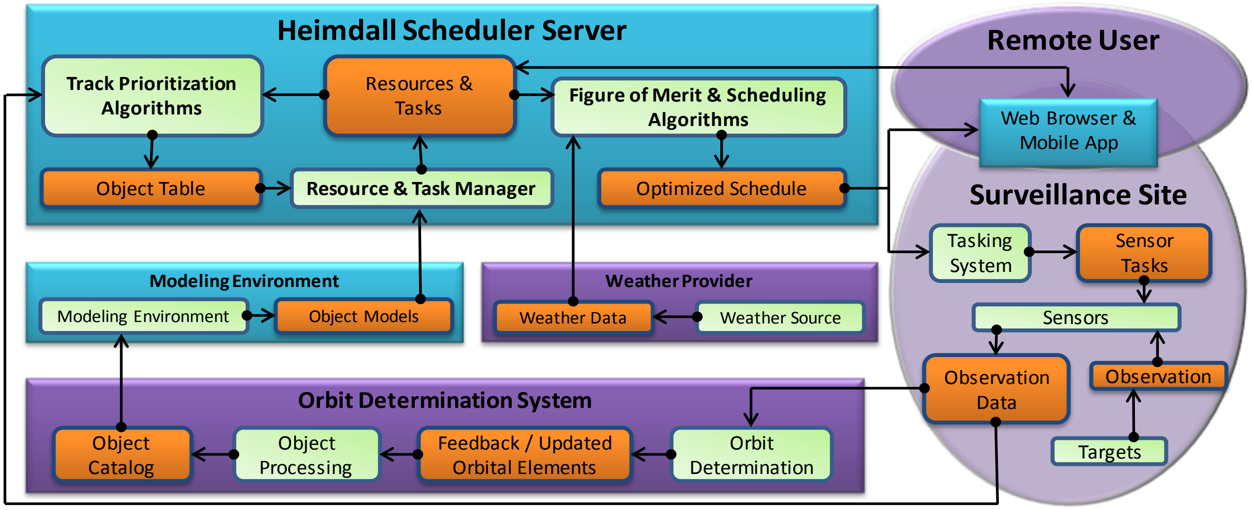

Fig. 1. Heimdall Architecture Diagram

While Heimdall is a centralized scheduling system, in the future it is possible that remote users from the Surveillance Site will have the ability to view the observation schedule via a web or mobile interface. The observation schedule is provided to the sensors at the Surveillance site which then collect data according to the schedule. An orbit determination system external to Heimdall processes the sensor data and updates the object catalog. The process then repeats using the updated information.

The operational process flow described here may be initiated on a regular schedule (e.g. daily) automatically or by an operator, or may be initiated automatically based on new events such as new object detection or unanticipated sensor downtime. The performance of the scheduling algorithms allows for responsive retasking of the observation schedule as often as needed.

## Baseline Model and Observation Schedule

A baseline model of the MSSS sensors with a target deck that included 836 orbiting objects (all unclassified operational satellites) was developed using STK and STK Scheduler. Initially STK's Import Wizard was used to import all operational unclassified satellites in the AGI object database. A utility was developed to create the sky search array and add 560 point targets to the STK scenario to represent it. The target array covered a +/- 56 degrees longitude range centered on MSSS site (the entire area visible to the sensors above the MSSS minimum sensor elevation angle) and +/- 10 degrees latitude centered on equatorial plane. The MSSS Facility was added using the STK object database. WFOV 1.2m, 1.6m, and Raven MSSS sensors were manually created and configured in STK that were then associated with the MSSS Facility. Separate sensors were created for day and night activity of the 1.6m sensor. A Gravity Well and two known geo belt Pinch Points were manually added to the scenario as well.

A one day schedule was created and STK Scheduler's COM interface with STK was used to ingest/create individual resources for each object, sky survey point target in the STK scenario, facility with associated sensors, gravity well, and pinch points. Objects were categorized into multiple groups including: High Priority GEO, GEO, MEO, and LEO. These groupings were used to differentiate overall priority of the objects within the groups. In the same manner, the sky survey points were grouped together too. STK Scheduler's template task functionality was used to generate tasks for each group. This resulted in the need to generate only 5 template tasks. Template tasks allow a user to generate multiple tasks that, except for 1 resource, are identical. Upon completion of a template task, a child task was created for each object within the group greatly reducing the amount of time required to add all of the objects and sky survey targets. Three more tasks were manually created for the gravity well and pinch point.

STK Scheduler's COTS algorithms were used to generate a baseline observation schedule, and were able to generate a valid, deconflicted observation schedule solution in less than 10 seconds.

**SSA Task and Resource Attributes**

As part of the Phase I effort, Orbit logic introduced new task and resource attributes specifically to support SSA scheduling. These new options, as well as the new SSA-specific FOM, are only visible to users when SSA features are enabled via STK Scheduler's Configuration interface (an option restricted to authorized customers).

New Resource attributes include:

- SSA Resource Type to specify a Sensor, Object Target, or Search Area Target
- For Sensors
    - Location (Lat/Lon/Altitude)
    - Cost provides for a cost rate or a single observation cost
- For Object Targets
    - Category (known vs. unknown)
    - Behavior (Passive, Active, Anomalous)
    - Covariance summary – The variance in the translation state represented as the semimajor axis uncertainty
    - Satellite Albedo (contributes to Visual Magnitude)
    - Satellite Radius (contributes to Visual Magnitude)
    - Last Observation Time
- For Search Areas
    - Predicted Target Density
    - Last Observation Time

New timeslot attributes include the following for specific observation windows of opportunity:

- Expected Information Gain - An information-theoretic metric used to quantifying the reduction in uncertainty over all targets within the FOV. Provided by CU Track Prioritization component.
- Observation Probability - The probability of a given target being within the FOV given a probability density function describing its translation state. Provided by CU Track Prioritization component.
- Range from Sensor to Target (contributes to Visual Magnitude)
- Solar Phase Angle between sun, sensor, and target (contributes to Visual Magnitude)

**SSA Figure-of-Merit (FOM)**

The FOM is used to score available sensor tasking opportunities - both before scheduling, and to compare candidate plans after scheduling. This loose coupling of the FOM and scheduling algorithms allows the FOM to influence algorithm decisions but not hamstring the algorithms by forcing specific assignment decisions.

When the Scheduling algorithms are run, several different de-conflicted solutions may be generated. Heimdall will use the configurable Figure-of-Merit (FOM) calculation to determine the most preferred schedule solution from the family of de-conflicted solutions. The FOM calculation features a series of terms, each designed to highlight a particular trait of a tasking opportunity or tasking assignment. The user can configure the FOM to specify weighting for each of the FOM terms. The solution with the greatest FOM value is the one that is presented to the user for review.

The FOM score for a schedule is computed as the sum of the score of each observation task included in the schedule. The score of an individual observation task is computed as the sum of the score of each of the following factors, with each factor multiplied by a user-configurable weighting value. Note that the Observation Task Score FOM is used to score each Timeslot (observation opportunity). The goodness of the overall schedule is determined based on the FOM score for each task assignment.

The Figure-of-Merit (FOM) score for an individual target (object or search area) equals the sum of each factor multiplied by its factor weighting value, which can any positive or negative integer value, or zero. The FOM factors for Object Tracking and for Search Areas are not identical, as one would expect. The FOM formula normalizes the Object Tracking and Search Area scores. The formula below represents the general FOM score computation for an individual object track or search area observation:

$$\text{Object (or Area) Score} = \sum_{Factor1}^{Factor\,n} FactorValue * FactorWeight \,/\, \text{(number of factors)}. \tag{1}$$

The general formula for the score of an overall plan, would therefore be as follows:

$$\text{Plan Score} = \text{TRACK} * \sum_{Incl.\ Object\ 1}^{Incl.\ Object\ n} ObjectScore\,/12 + \text{SEARCH} * \sum_{Obs.\ Area\ 1}^{Obs.\ Area\ n} AreaScore\,/7. \tag{2}$$

The divisors for the Object and AreaSearch scores are the number of non-zero factors.

We cannot identify the specific factors in our SSA-specific FOM in this paper, but the multi-factor SSA-specific FOM developed includes factors associated with orbit knowledge, priorities, and the value of a specific observation to improving overall knowledge. Note that the FOM can easily be modified to add new factors or change the weighting of various factors without requiring a rewrite of scheduling algorithms.

### Track Prioritization Component

CU has developed tools based on Finite Set Statistics (FISST) to aid in prioritization of new and previously known tracks. This includes the quantification of information gain, uncertainty propagation, and target density in the sensor field of view. The FISST-derived filter used in Heimdall is the Cardinalized Probability Hypothesis Density (CPHD) filter, illustrated conceptually in **Fig. 2**.

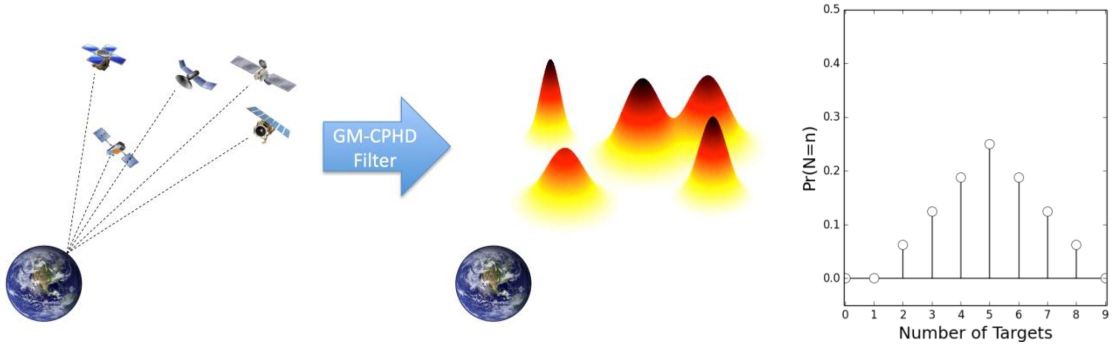

**Fig. 2. Illustration of CPHD Filter Components**

Given a set of measurements at a prescribed time, the CPHD filter predicts and corrects an intensity function describing the spatial distribution of targets, as well as a cardinality distribution providing the probability that a given number of targets are present. The update equations for the intensity and cardinality are coupled, meaning that any measurements collected provide information for both.

The expected information gain is computed using the Rényi divergence (Ristic, et al. (2011))

$$R(\boldsymbol{u}) = \frac{1}{\alpha - 1} \log \sum_{n=0}^{\infty} p_1(n;\boldsymbol{u})^{\alpha} p_0(n)^{1-\alpha} \cdot \left[ \int s_1(\boldsymbol{x};\boldsymbol{u})^{\alpha} s_0(\boldsymbol{x})^{1-\alpha}\, d\boldsymbol{x} \right]^n \tag{5}$$

and, when used with the CPHD, assumes objects are independent and identically distributed (I.I.D.). The metric allows quantification of the difference between a prior and posterior multitarget PDF based on a proposed sensor tasking vector $\boldsymbol{u}$. The multitarget PDF is parameterized by the cardinality distribution, $p(n)$, and the spatial density of targets, $s(\boldsymbol{x})$. The subscripts 0 and 1 represent the prior and posterior respectively. We use a Gaussian mixture approximation of the target density, i.e.,

$$s(\boldsymbol{x}) = \sum_{j=1}^{J} w_j p_g\left(\boldsymbol{x}; \boldsymbol{m}_j, \boldsymbol{P}_j\right) \tag{5}$$

where $\boldsymbol{m}_j$ and $\boldsymbol{P}_j$ are the mean and covariance matrix for the *i*-th component. Upon selecting the tuning parameter $\alpha = 0.5$, the reward function is then

$$R(\boldsymbol{u}) \approx -2\log \sum_{n=0}^{\infty} \left(\frac{p_1(n;\boldsymbol{u})}{N_1^n}\right)^{1/2} \left(\frac{p_0(n)}{N_0^n}\right)^{1/2} \cdot \left[\int \left(\sum_{i=1}^{J_0}\sum_{j=0}^{J_1} w_i w_j K_{i,j} p_g\left(\boldsymbol{x}; \boldsymbol{m}_{i,j}, \boldsymbol{P}_{i,j}\right)\right) d\boldsymbol{x}\right]^n \tag{6}$$

$$K_{i,j} = \frac{1}{\left|2\pi(\boldsymbol{P}_i + \boldsymbol{P}_j)\right|^{1/2}} \exp\left[-\frac{1}{2}(\boldsymbol{m}_i - \boldsymbol{m}_j)^T(\boldsymbol{P}_i + \boldsymbol{P}_j)^{-1}(\boldsymbol{m}_i - \boldsymbol{m}_j)\right] \tag{7}$$

$$\boldsymbol{P}_{i,j} = \left(\boldsymbol{P}_i^{-1} + \boldsymbol{P}_j^{-1}\right)^{-1} \tag{8}$$

$$\boldsymbol{m}_{i,j} = \boldsymbol{P}_{i,j}\left(\boldsymbol{P}_i^{-1}\boldsymbol{m}_i + \boldsymbol{P}_j^{-1}\boldsymbol{m}_j\right) \tag{9}$$

This information gain metric produces a zero value in the case where the prior and posterior PDFs are the same, and increases as the difference between the PDFs increases. Component covariance updates depend on the previous uncertainty, observation geometry, and measurement quality, and may be better for certain regions at certain times of the day. The full multitarget information gain also depends on the number of objects in the field of view, and will generally be higher if more objects are observed. Using the information gain metric as a selector in sensor tasking allows the user to maximize the efficiency of measurements taken in driving down the uncertainty associated with the multitarget problem.

The current covariance output for each object is the uncertainty in semimajor axis, computed by converting the Cartesian covariance matrix to orbit element space using the unscented transform. The approach allows a single number to represent the uncertainty for each object and gives the user an option to focus on collecting measurements for objects with the highest uncertainties.

For search areas, CU's tools integrate the estimated intensity function $s(\boldsymbol{x})$ over the sensor field of view to provide the expected number of objects visible and target density. As mentioned above, the intensity function is maintained as part of the CPHD filter. The metric allows for careful selection of search areas in sensor tasking, as dense regions are more likely to produce new targets for tracking.

To enable the follow-up tracking of new targets, CU has extended current tools based on the constrained admissible region (CAR) for use in Heimdall. Such techniques assume that more than one observation of a new target is available in the initial track that identified its possibility. CU has implemented tools to convert such tracks into angle and angle-rate measurements.

Given a four-parameter measurement set (angles and rates), it is possible to constrain two additional states to produce a six-parameter initial orbit description. CU has implemented a set of constraints in eccentricity and semimajor axis, designed to look for GEO-crossing objects. The Fig. 3 below illustrates the concept of the CAR. On the left, a set of constraints in eccentricity and semimajor axis have been mapped into range/range-rate space. The region of overlap, the CAR, represents all possible solutions matching the original measurement set and fitting the constraints. The region can be discretized into a set of possible solutions, represented by a weighted sum of Gaussian distributions, as indicated on the right. Filtering subsequent measurements allows refinement of the solution for the new object.

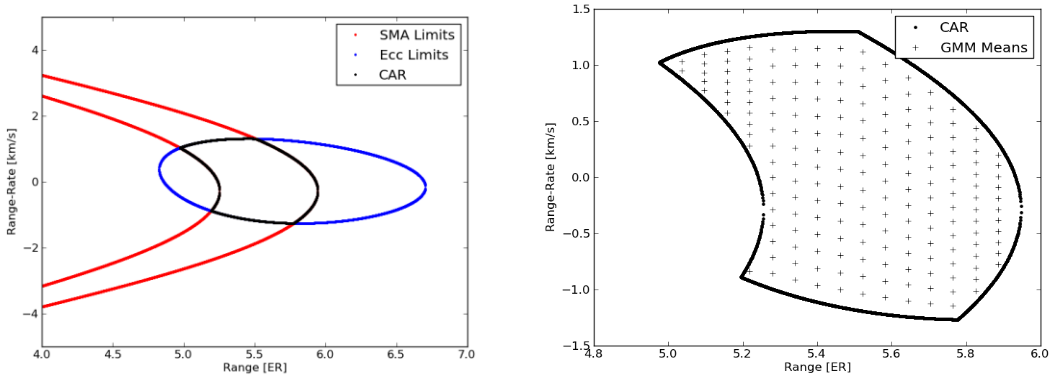

**Fig. 3. Representation of the CAR**

To demonstrate the follow-up observation of new targets within Heimdall, CU has developed a test case using the baseline scenario that includes a simulated new target. CU and Orbit Logic defined the method for the track prioritization software to notify the scheduler that a new target is available. Upon identifying the new object and initiating a CAR-based IOD solution, this allows for later demonstration of follow-up tracking.

In addition to the previously mentioned capabilities, CU has integrated software tools that enhance the Track Prioritization component software. The CU-TurboProp package provides C/C++-code for fast propagation of satellite orbits and allows for non-Gaussian uncertainty propagation using the AEGIS Gaussian mixtures algorithm. Additionally, the CU-TurboProp package provides a suite of models to simulate perturbing forces, due to the non-uniform gravity field, drag, SRP, and 3rd body effects. The added Earth Orientation Parameter (EOP) package allows for accurate coordinate system reduction based on the latest International Astronomical Union (IAU) standards.

CU has implemented a series of tests to demonstrate the computation of values for the FOM, including information gain and target density. As an initial check, information gain is computed for a small sample of objects over a one-day period. A total of five objects from the public TLE catalog are simulated. Two are typically within the same FOV and the other three are separated such that only one target is visible for a given sensor task. Of the lone objects, one is initialized with a large radial uncertainty, one with a large in-track uncertainty, and the final with a large cross-track uncertainty. As seen in Fig. 4, the information gain for the two-object case is the largest throughout the day. The cross-track uncertainty case is the largest of the single objects, which is expected because the angle measurements provide the most information in the cross-track and in-track directions. The radial and in-track cases alternately produce more information gain than one another due to the orientation of the uncertainty ellipsoid, which is aligned in the orbital plane and rotates from radial to in-track.

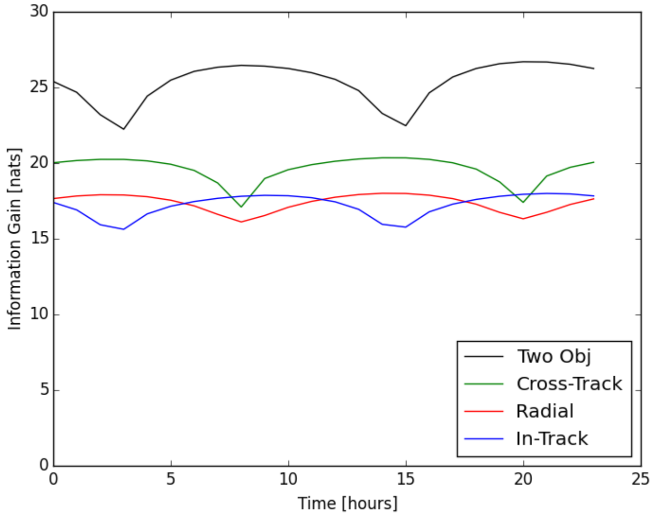

Fig. 4. Information Gain over 1 Day

To demonstrate the target density function, a sample object catalog is generated, in which the initial state of all objects, as determined from STK, is propagated back three days. A set of sparse angle measurements are generated, in which objects are observed for 10 minute arcs, followed by 12 hour gaps. A filter is initiated for each object, beginning with a diagonal covariance matrix, and processing the measurements until the time UTC Sept 19, 2014 00:00:00. This produced an object catalog with realistic covariances for use in computing the expected target density as described above.

Fig. 5 provides the results of the target density test, which is computed by tasking the sensor to point at all search areas at a single point in time.

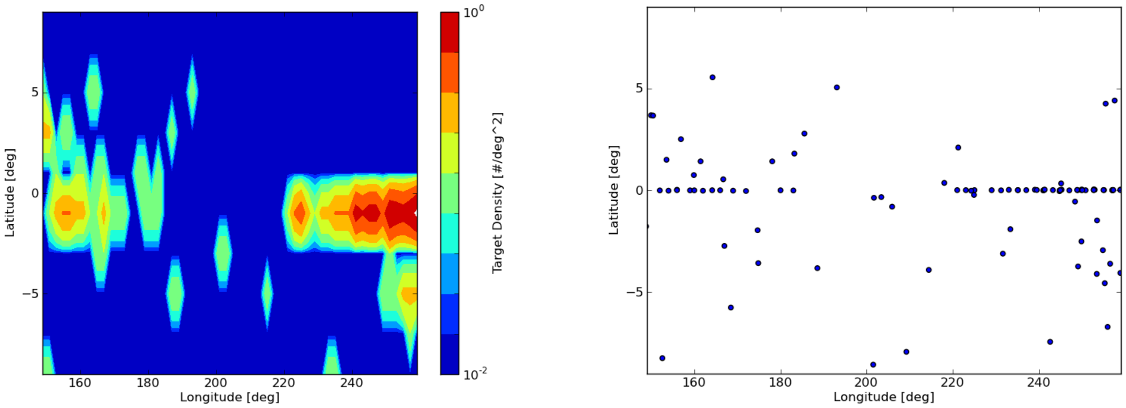

Fig. 5. Target Density at 00:00:00

On the left, the computed target density is plotted versus the longitude and latitude of the defined search area regions. On the right, the actual locations of all visible objects are plotted. As expected, certain regions exhibit a much higher density of targets than others, making them more desirable. The apparent 1 degree shift in latitude between the actual targets and predicted density is due to the discrete search area bins, which are assigned at 2 degree intervals. All objects with latitude between 0 and -2 degrees are therefore grouped in the -1 degree latitude bins. A number of search areas, particularly away from the equator, contain no known targets.

## Responsive Retasking

In order to provide responsive retasking, Heimdall includes a file dropbox to allow scheduler to process new object targets for schedule tasking. The new object targets are defined in text files with an ID and Two Line Element (TLE) data. Scheduler detects new object files delivered by the CU Track Prioritization component and automatically updates any remaining portion of the current operational schedule to consider a new observation task(s) for the new object.

## 3. HEIMDALL DEMO AND RESULTS

The Heimdall proof-of-concept was demonstrated for MSSS SSA sensor tasking for a 24 hour period to attempt observations of all operational satellites in the unclassified NORAD catalog, observe a small set of high priority GEO targets every 30 minutes, make a sky survey of the GEO belt region accessible to MSSS sensors, and observe particular GEO regions that have a high probability of finding new objects with any excess sensor time.
This Heimdall prototype software paves the way for further R&D that will integrate this technology into the MSSS systems for operational scheduling, improve the software's scalability, and further tune and enhance schedule optimization.

### Baseline Results

Using the baseline developed as stated in the previous section, multiple algorithm runs were made using different algorithms, and a hybrid run with a minor manual update by an experienced operator. Results for all runs are shown in Table 1. The Baseline STK scenario is depicted in Fig. 6. The "best" deconflicted baseline observation schedule was generated by STK Scheduler's COTS algorithms in less than 4 seconds with a minor manual update prior to running the algorithms. This solution was generated prior to any Heimdall enhancements being implemented. That schedule was then animated in STK (Fig. 2). Baseline software performance bodes well for scalability of the system. The baseline schedule will be used for comparison metrics against the Heimdall SSA enhanced software.

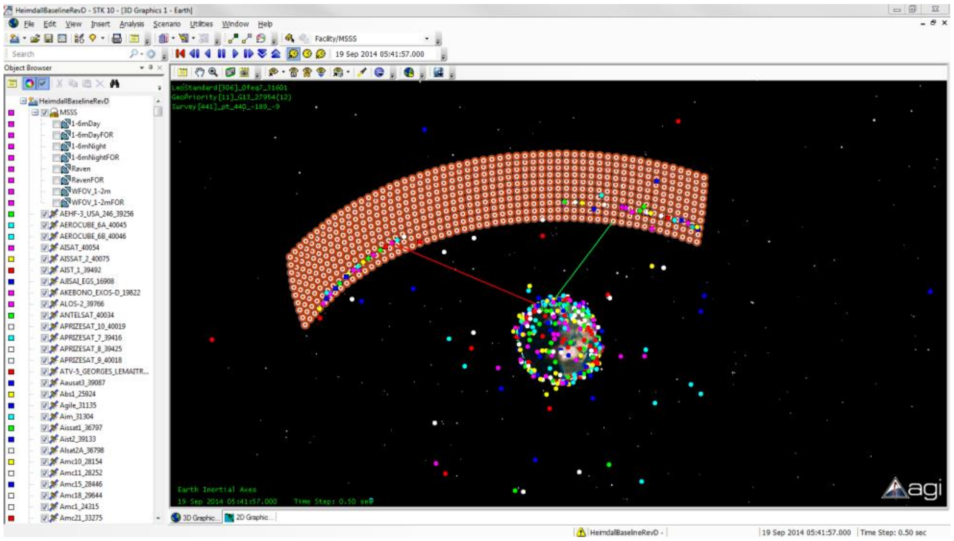

**Fig. 6. Animation of Baseline Schedule in STK Model**

1. Algorithms took less than 10 seconds to generate a valid schedule solution
2. 2040 observations scheduled (over 24 hours)
3. All satellites requested to be observed once per day
4. 95.6% of possible observations scheduled (best solution)
5. All of the unassigned observations were targets without any access or LEO targets conflicting with higher priority observations . . . which is consistent with the priorities assigned (GEO > MEO > LEO) for this demo run (priorities may be changed at any time by an operator)
6. Full survey region covered (at night within lighting constraint)
7. All High Priority GEO Targets were observed every 30 minutes both day and night, except when in shadow
8. Gravity Well and Pinch Point areas were observed for long durations (>1hr each)

**Table 1. Baseline Schedule Algorithm Results**

| Algorithm | Total Assigned | % Assigned | Total Assigned Dur (sec) | Schedule Score | Time to Deconflict (sec) | Gravity Well Dur | Pinch Point 1 Dur | Pinch Point 2 Dur |
|---|---|---|---|---|---|---|---|---|
| OnePass | 1969 | 92.27% | 69986.696 | 715848.881 | 3.463 | 04:56:00 | 03:34:47 | 00:00:00 |
| Sequential | 1959 | 91.80% | 69786.696 | 714497.037 | 3.744 | 04:56:00 | 03:34:47 | 00:00:00 |
| Multi-Pass | 1969 | 92.27% | 69986.696 | 715848.881 | 8.596 | 04:56:00 | 03:34:47 | 00:00:00 |
| Random | 2040 | 95.60% | 48419.812 | 725905.437 | 12.714* | 01:48:00 | 00:10:00 | 00:10:00 |
| Hybrid | 2041 | 95.64% | 71406.543 | 725557.804 | 11.544 | 04:56:00 | 02:12:32 | 01:22:15 |

### Heimdall Prototype Software Demonstration

The Heimdall software was run for the customer to demonstrate key features of the prototype software. The remote demonstration was run on a standard laptop with a connection to the customer via standard web meeting software. The SSA-Enhanced Baseline scheduling scenario was the focus of the demonstration, with SSA-specific attributes explored and executed, and schedules inspected. Both general SSA sensor scheduling as well as responsive retasking was demonstrated. At a high level, the following steps formed the Heimdall demonstration:

- Review of STK scenario system model (sensors and RSO and search area target objects)
- Review of scheduling resource definitions & attributes
- Review of observation task definitions & attributes
- Review of Task Prioritization results
- Demo of Observation scheduling algorithm run
- Review of scheduling result in Gantt and Task & Resource Editors
- Review of schedule in 3D animation
- Analyze an unscheduled observation, update its score manually, re-run scheduling algorithms, assess results
- Rerun Track Prioritization script on the modified task and rerun scheduling algorithms to reset the schedule
- Demo New Target notification and automated responsive observation schedule updates
- Review of selected schedule reports

The demonstration was based on the same scheduling scenario used throughout the Heimdall effort and incudes three MSSS sensors, all unclassified operational satellites (836) in various orbits, and search area targets.

### Heimdall FOM Testing Results

To validate the Figure-of-Merit, each FOM factor was isolated and confirmed. To do this, a small schedule that only spanned 2 hours was generated with 2 objects and 2 survey targets. A Microsoft Excel spreadsheet was created to independently calculate the schedule score. In the spreadsheet, the weighting factors were all set to zero except for the Global Priority Weighting and 1 other factor which were both set to “1”. The same was done on the FOM configuration UI in Heimdall and the One-Pass algorithm was re-run. The Excel result was compared to the Heimdall schedule score to validate the accuracy of the calculations. This was repeated for each FOM factor which was noted. After all factors had been validated, the scores as determined by Heimdall, were summed into 2 groups, the object factors and survey factors, and divided by their respective number of factors. These 2 numbers were then added together to be used to validate against a multiple FOM factor scoring run. At this point, all FOM factors were set to “1” in the FOM configuration UI and the One-Pass algorithm was re-run. This total score as output by Heimdall was then validated against the manually calculated total score.

It should be noted that while the FOM score validation was successful from a computational perspective, additional work needs to be done to determine the proper weighting balance between the various FOM factors to achieve the best possible catalog knowledge with available observation assets.

### Heimdall Algorithm Testing Results

Performance runs were executed with multiple algorithms to determine the “best” solution. The quality of the schedule was measured by the percentage of assigned tasks, an assessment of the score/priority of unassigned tasks, and an assessment of the duration of search area observation dwell times. Table 2 contains metrics for Baseline and sample SSA-Enhanced runs.

**Table 2. Baseline and SSA Algorithm Results**

| Algorithm | Type | Total Assigned | % Assigned | Total Assigned Dur (sec) | Schedule Score | Time to Deconflict (sec) | Gravity Well Dur | Pinch Point 1 Dur | Pinch Point 2 Dur |
|---|---|---|---|---|---|---|---|---|---|
| One-Pass | Baseline | 1969 | 92.27% | 69986.696 | 715848.881 | 3.463 | 00:04:56 | 03:34:47 | 00:00:00 |
| Sequential | Baseline | 1959 | 91.80% | 69786.696 | 714497.037 | 3.744 | 00:04:56 | 03:34:47 | 00:00:00 |
| Multi-Pass | Baseline | 1969 | 92.27% | 69986.696 | 715848.881 | 8.596 | 00:04:56 | 03:34:47 | 00:00:00 |
| Random | Baseline | 2040 | 95.60% | 48419.812 | 725905.437 | 12.714 | 01:48:00 | 00:10:00 | 00:10:00 |
| Hybrid | Baseline w/lock | 2041 | 95.64% | 71406.543 | 725557.804 | 11.544 | 00:04:56 | 02:12:32 | 01:22:15 |
| One-Pass | SSA | 2059 | 96.49% | 61638.466 | 121990180.96 | 4.941 | 03:39:00 | 02:02:00 | 00:00:00 |
| One-Pass | SSA w/lock | 2061 | 96.58% | 66593.422 | 122194102.16 | 3.695 | 03:39:00 | 02:02:00 | 01:22:00 |

Results Notes:

1) The Random algorithm run was done as “best of 100”.
2) The One-Pass SSA run “w/lock” refers to the Pinch Point 2 task being manually assigned and locked for the entire timeslot duration prior to running algorithms.

Comparing the new SSA-enhanced One-Pass solution against the original non-SSA-enhanced One-Pass solution, showed that

1) The number of included tasks increased by more than 4% (+90).
2) More time was spent observing the Gravity well and Pinch points.

It should be noted that while the improvements above are encouraging, comparing the pre-enhanced run to the SSA-enhanced run is to some extent comparing apples to oranges. The pre-enhanced schedule solution is completely blind to SSA-specific factors. Probably the most significant result is that we can influence the scheduling algorithms to be driven by SSA-specific FOM scoring fed by computed SSA-specific factors, as is detailed below.

## 4. HEIMDALL CONCLUSIONS AND NEXT STEPS

The Heimdall software for SSA sensor tasking provides greatly improved performance over manual tasking, improved coordinated sensor usage, and tasking schedules driven by catalog improvement goals (improved orbit knowledge, etc.). The improved performance also enables more responsive sensor tasking to address external events, newly detected objects, newly detected object activity, and sensor anomalies. Instead of having to wait until the next day’s scheduling phase, events can be addressed with new tasking schedules immediately (within seconds or minutes).

Perhaps the most important benefit is improved SSA based on an overall improvement to the quality of the space catalog. By driving sensor tasking and scheduling based on appropriate relevant, computed factors, better decisions are made in the application of available sensor resources, leading to an improved catalog and better information about the objects of most interest.

The Orbit Logic/CU team are highly encouraged by the results of the Heimdall demo effort. Based on the results presented in this report, the following technical conclusions can be easily drawn:

1. The COTS software Baseline solution provides a good method to quickly generate valid, de-conflicted MSSS observation schedules
2. The prototype Track Prioritization component shows that a tractable and scalable FISST-inspired SSA observation opportunity solution is achievable
3. The prototype Figure-of-Merit (FOM) and Scheduler algorithms show that a configurable SSA-specific FOM can be applied for optimized SSA observation scheduling
4. The prototype Heimdall software shows that routine, scalable SSA observation schedule optimization is achievable and has the performance to support responsive re-tasking based on new events such as new detected objects or sensor maintenance issues

## Future Work

The results presented in this paper were achieved in the initial Phase I SBIR effort sponsored by the Air Force Research Laboratory (AFRL). Future work (under a Phase II SBIR contract expected to be awarded by the time of the 2015 AMOS Conference, and other expected contract awards) includes:

- Integration of the prototype Heimdall software with an operational observation system
- Comparison of Heimdall observation schedules against current operational observation schedules
- Refinement of SSA FOM factor computations
- Investigation of various factor weighting combinations to determine the best settings for optimal catalog maintenance
- Performance and scalability improvements for larger space object catalogs

## References

**Ristic, Branko, Ba-Ngu Vo, and Daniel Clark, "A Note on the Reward Function for PHD Filters with Sensor Control",** ***IEEE Transactions on Aerospace and Electronic Systems*****, Vol. 47, No. 2, 2011, pp. 1521-1529.**