Almost-Optimal Sensor Tasking using Auction Methods

Richard S. Hujsak
Analytical Graphics, Inc.

ABSTRACT

The problem for efficient use of radar assets for space surveillance is unsolved. Partial solutions have been developed, such as covariance based tasking algorithms, but not all operational requirements for tracking support are quantified by accuracy, for example monitoring a non-cooperative maneuvering satellite for purposes of conjunction assessment. The following effort exposes a different approach, where competing requirements, each with quite different metrics, can compete for sensor resources and an efficient decision making process will provide for economic use of radar resources and thereby allowing greater use of the radar sites for other missions. The Auction we employ here is somewhat naïve, as compared to sophisticated methods employed by domain experts, and yet provides evidence that the method can be quite useful for making assignments where competing requirements can stress the tracking systems.

1. OVERVIEW

The modern space catalog problem includes tens of thousands of large objects, some of which are active while most are inoperative. The problem is to efficiently task various tracking devices (radar and optical) to collect tracking data in a way that is efficient and supports good accuracy for all orbits. The approach taken here is to determine the optimal sensor coverage for each high interest object and to assign a "value" to each "track" of data. Typically the "value" represents the expected improvement in orbit accuracy as a result of using that "track". In a marketing sense each object is "acting in its own best interest" or is its own "agent". The problem then is to sort through all of the satellites competing for radar resources and to make the best possible assignments of tracking so that the overall performance of the entire system is "almost optimal"[1].

The problem is complicated by the variety of sensors available to track satellites. For example an optical instrument with a limited field of view can track a very small number of targets at any one time, and there is a non-zero loss of time for "slew and settle" between tracking opportunities. Conversely, phased-array radar sites can track a fairly large number of targets in near-real-time, although there is some loss of accuracy if the target set gets to be very large due to shorter track lengths as the radar time-shares resources.

We constructed a modest simulation of a subset of the Space Catalog, using three hundred objects from the general catalog and arbitrarily assigned all objects to be "active" and "high priority". These objects are representative of a high priority subset of the total catalog, which normally demand a disproportional amount of radar resources. Our objective is to find an efficient level of tasking for this subset so that the radar resources can be better used to track the rest of the catalog.

Having constructed the testbed, we set about to find a method to make the assignments that would yield an efficient use of the sensor resources and maintain the limits of each of the sensors in terms of number of objects under track concurrently, while meeting the notional mission priorities. The first approach we took was to establish a "Rule of Thumb" Tasking algorithm based on orbit analyst intuition (geometric distribution and time distribution of tracking). The second was to introduce concepts from econometrics, with each satellite computing a "value" for each requested track, a value that can be used to construct a "bid" for the resources. Both methods were compared to a baseline where all possible tracking from all possible radars was used to determine the "best" solution. Results indicate that there are significant differences in performance between different Tasking algorithms, with the Auction Method being better than the "Rule of Thumb" method.

2. SIMULATED TESTBED

---

[1] "Optimal" strictly means the "best" solution at every point. Finding such a solution is "NP-hard", which means that it is computationally prohibitive. Therefore, an "almost optimal" solution is sought which is computationally tractable.

The simulation testbed is constructed from the current space catalog. The objects are all LEO objects, and are taken from the long-dead section of the catalog, beginning with object 22. The first 300 payloads in the catalog that are in LEO circular orbits were selected and treated as a representative set of orbits. Of these 300, we postulate that for 290 objects accuracy is the driving requirement. For the remaining 10 satellites we assume frequent maneuvering is the issue driving the Tasking Algorithm. The current set of phased array radar sites were selected as a representative tracking system, and were modeled to include radar face orientation (approximately) and Field of Regard (F.O.R.). A candidate list of all possible tracking opportunities was constructed by running "Access" in Satellite Tool Kit (STK), where tracking for each radar face is limited to the proper F.O.R. The aggregate set of all possible radar tracks for all satellites comprises the search domain for an almost-optimal tasking algorithm.

3. GENERATING A "BEST PERFORMANCE" BASELINE

The best possible performance, in terms of accuracy, is achieved if we use all possible radar tracks from all radar faces. Any subset of tracking data will then suffer some degradation from this baseline, and that degradation will be the measure of effectiveness of the tasking algorithm.

For the purposes of establishing the best possible performance we generated radar tracking data for all possible tracks and performed orbit determination using Orbit Determination Tool Kit (ODTK) and archived the ephemeris and covariance for each of the 300 objects. We also archived the "Truth" ephemerides from the simulations to be used for performance analyses. Radar performance statistics were postulated from theoretical performance for UHF radar sites and with some consultation with radar analysts. Each phased array radar site was assigned range white noise sigmas of 25 m and angular white noise sigmas of 0.005 degrees. In every case we ignored the range rate data. Radar reporting was assumed to be one observation every 10 seconds throughout the F.O.R. Observations were formatted in the typical Space Surveillance Network (SSN) "B3" observation format.

Fig 1 provides the aggregate accuracy (2-sigma) from processing all 300 satellites, given all possible tracks from all possible radar sites. Some of the satellites exhibit rather poor orbit error. These are low inclination satellites which enjoy limited tracking. There are four such satellites and they will be ignored in the final analysis, as they distract from the experiment for Auctions. The net testbed for Auction methods will be 296 satellites.

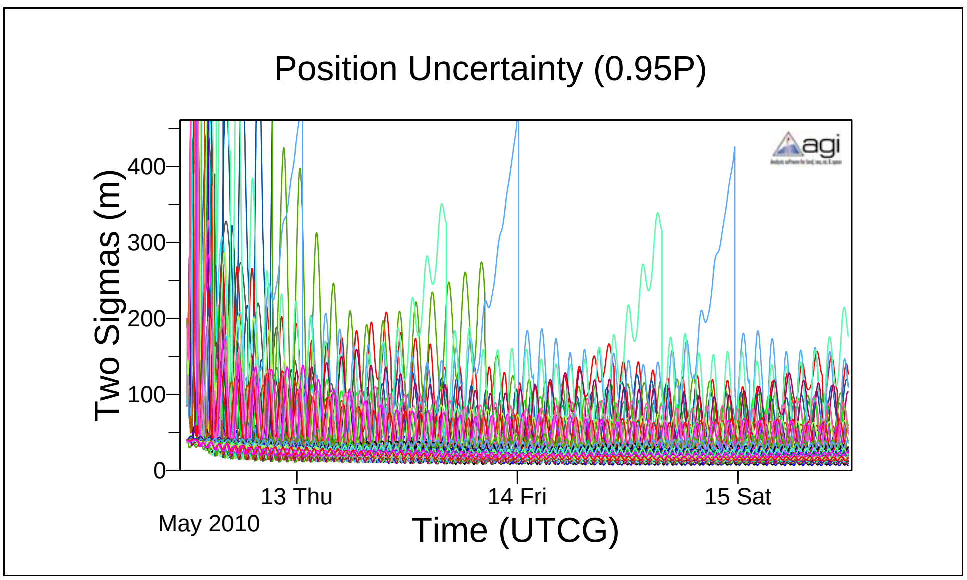

Fig. 1: 2-Sigma Accuracy for All Possible Tracks

The number of observations created this way is enormous (2.6 million), and would be difficult to process in current space surveillance centers. But the point of this exercise is to establish a "best" performance metric. Curiously enough, for 300 satellites the loading on each facility falls well below our arbitrary limit of 30 concurrent tracks.

However the maximum of 30 concurrent tracks would be easily exceeded if the test case was for 1000 high interest targets.

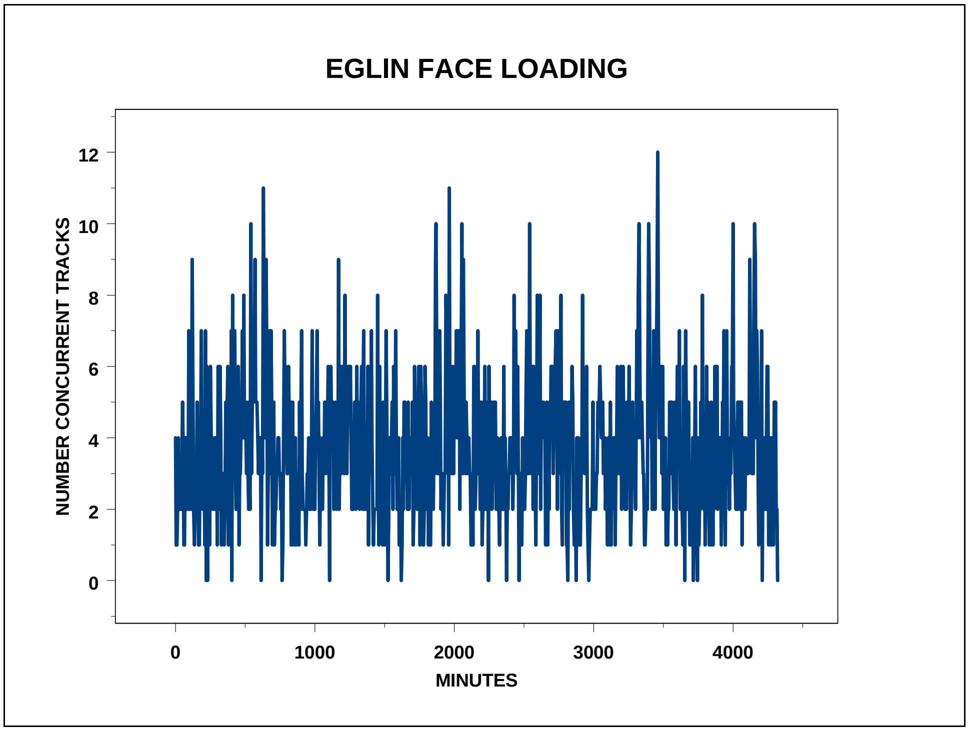

Fig. 2: Concurrent Tracks at Eglin, "All" Method

4. A "RULE OF THUMB" TASKING ALGORITM

A "Rule of Thumb" tasking algorithm is simply that, an aggregation of analyst's intuitive rules assembled in software and called a "tasking algorithm". This approach emulates the historical development of tasking algorithms for space surveillance, which have been of questionable utility. These rules are as follows:

1. Repeated tracking by the same radar site does not necessarily improve the orbit, unless the tracks are separated by 12 hours.
2. Successive tracks from different radar sites which provide the greatest geometric coverage of the orbit ellipse will provide the best accuracy.
3. Short tracks are sufficient[2], in this case = 1 minute duration.
4. Each radar face has a maximum number of targets that can be tracked simultaneously; we will use 30 as a representative value[3].
5. If one face of a radar site is used to track an object, there is no advantage to continuing the track with another face of the same radar site[4].

The rules were implemented, tracking data simulated, and orbit determination with ODTK performed. The filter output ephemerides for each of the 296 satellites were saved for comparison to the simulated Truth and to accumulate performance statistics. The "Rule of Thumb" tasking algorithm reduced the number of observations from 2.6 million to 100 thousand. There is an impact on accuracy, as discussed in Section 6.

---

[2] This is not the recommendation of the author, but seems to be the pervasive "wisdom" within AFSPC.

[3] This limitation is simply a matter of power. The actual limitations for each radar site are not known to the author. The value of 30 simultaneous objects is purely an invention for purposes of this analysis.

[4] However if the target is maneuvering during the track there is substantial information that can be derived about that satellite and its capabilities by continuing the track.

5. AN AUCTION ALGORITHM

According to Shoham and Leyton-Brown [1], auctions address problems of “allocating (discrete) resources among selfish agents in a multi-agent system”.

Our Auction algorithm consists of agents (e.g. satellites) which engage in a competition for resources (e.g. radar tracking). The mechanism for the competition is a bidding process, where the stake-holders compete for the resources through a bidding process. A decision-making process attempts to find the best combination of “bids”, resulting in the “best” allocation of resources. This description is far simpler than the algorithm. The advantages to an Auction are that the agents do not need to have the same reasons for making their bids; each can have different reasons. Therefore an agent that maneuvers frequently (where maneuver detection is the driving motivation) can compete directly with an agent that is used for radar calibration (where accuracy is paramount) and both can compete with a reentering agent (where the orbit changes rapidly due to drag).

In this case we need two types of bids, one type for the agents that maneuver frequently and another for the remaining agents that require refined orbit accuracy. The “bids” are independent of either metric, being unitless. The decision process will seek to maximize the value to the radar system. The process of optimizing 296 agents across 13 radar “faces” is NP-hard, which means that the search for the absolute “best” (or optimal) solution is computationally prohibitive, and may not be solvable in a realistic timeline. That drives the decision-making process to use approximating methods and will generally create a suboptimal solution; how suboptimal is a matter of analysis.

The trick to making the auction work for agents with diverse interests is in determining an appropriate set of rules for constructing each bid, which is unitless, such that the bids relate to the relative importance or operational priority for each agent. Sophisticated bidding techniques are discussed in the literature [3] however we prefer a simple bidding strategy for a first naïve auction. To that end we allowed each agent to make a “natural” bid (accuracy requirements yield a bid based on predicted orbit error, maneuvering agents bid the revisit frequency, etc). The (simple) bidding method employed here changes the bid linearly based on the orbit error (for accuracy agents) or on the length of the preceding tracking data gap (for maneuvering satellites). The graphs we used are given in Fig. 3.

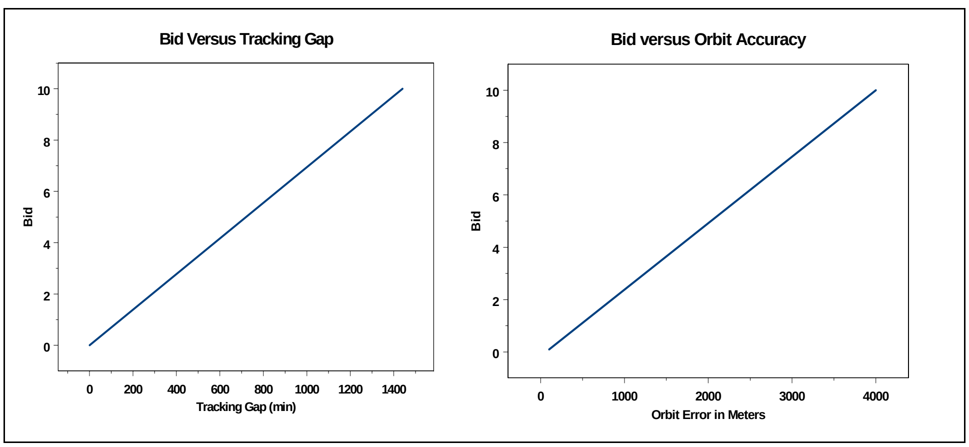

Fig 3: Simple Bidding Strategies for Maneuvers and Accuracy

The actual Auction process proceeds in steps of 3 hours across the three days, and Tasking is scheduled for a single step before proceeding to the next step. Then the results of the historical schedule are used to establish the historical performance and motivate the agents to make an appropriate bid based on that performance. For example, planned tracking events are processed to reduce the covariance and accuracy agents will bid appropriately in the next computational segment. Tracking events to monitor maneuvering satellites may either reduce the need for more tracking or may motivate acquiring more tracks at a higher bid.

The usual Auction method selects bids from the ensemble of agents for the resources (tracks) and then resumes bidding for remaining resources. We have taken an approach in simply taking the highest bids. In this way we have sped up the bidding process although it might be considered to be naïve.

For purposes of the Auction algorithm we have assumed that 3 minutes of track will be "purchased" if it is available.

6. ANALYSIS OF PERFORMANCE

Graphs like Fig 1 are nearly impossible to decipher and compare, so we extracted the performance of each method one satellite at a time, and then attempted to accumulate the performance statistics for comparison purposes. The performance statistics we generated were the mean orbit error and standard deviation orbit error on a uniform time grid of one minute, extending over the 3 days of analysis. The results are presented as a scalar RMS of the standard deviations. The "best" possible accuracy is achieved with all tracking data, regardless of the rules of practicality and is given in Fig 5. The peculiar scalloped accuracy is due to gaps in tracking coverage (see Fig. 6).

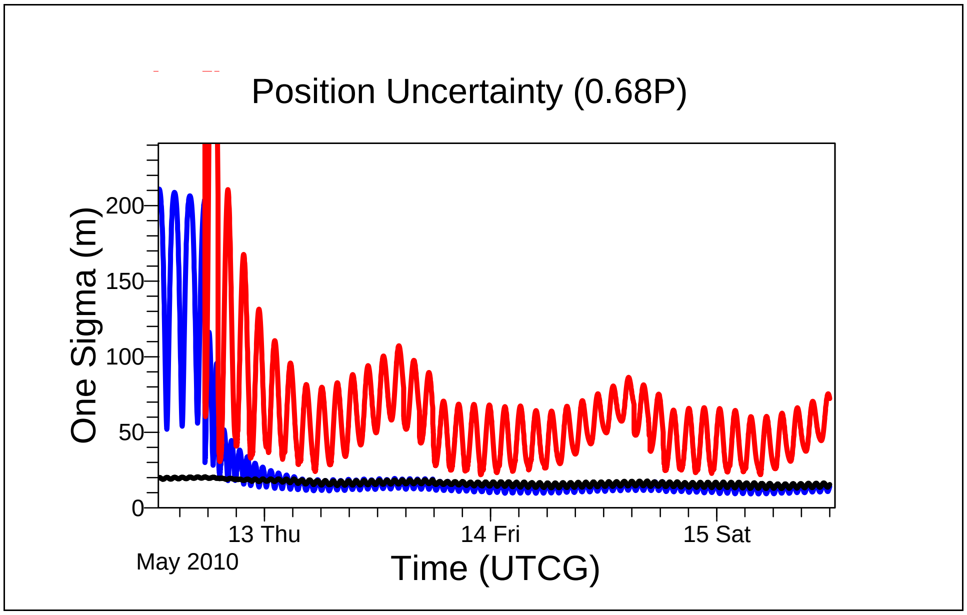

Fig. 5: Sample Performance of "All" Tracks (Optimal)

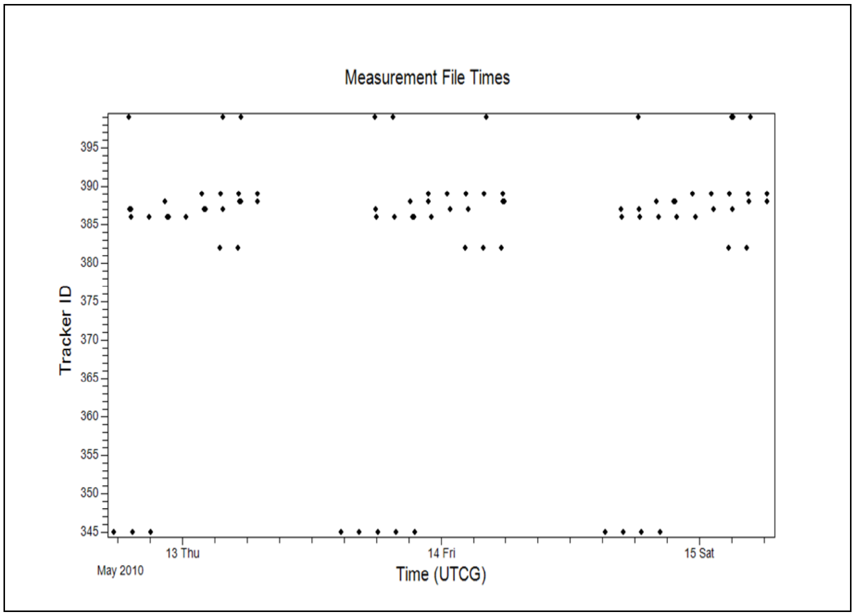

Fig 6. Gaps in Tracking Schedule

For convenience the accuracy of any single instance of a satellite can be represented as a scalar RSS of the three orbit error components, as in Fig. 7. Note the dashed line which indicates a converged accuracy of slightly more than 60 meters (RSS of one-sigma values).

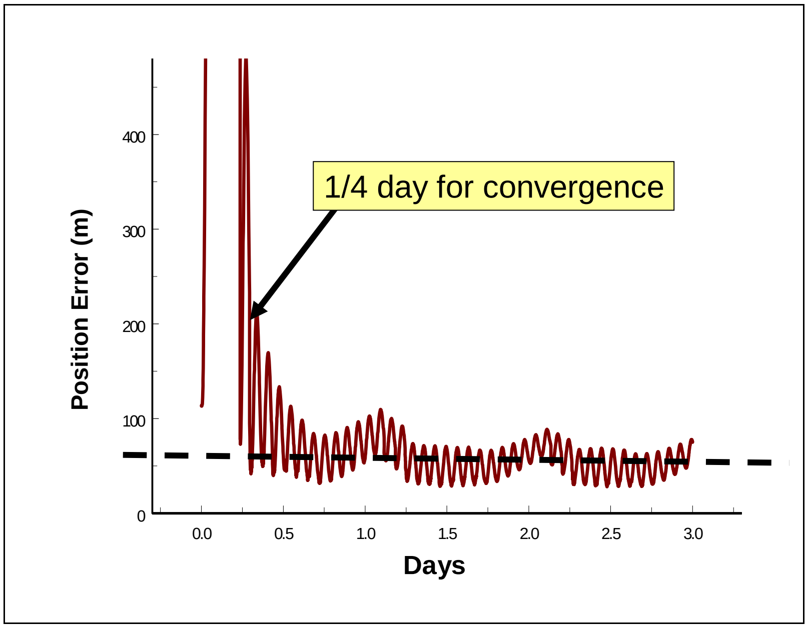

Fig 7: RSS of Position Error Components, Satellite 29, All Tracks

The corresponding graph of RSS values for the "Rule of Thumb" Tasking method is given in Fig 8. The dashed line gives an approximate overall accuracy of nearly 100 m that of the optimal solution and suffers a longer convergence time. The Auction tasking method gives better performance, as illustrated in Fig 9. The benefit of a 3 minute track duration versus the Auction itself has not been resolved.

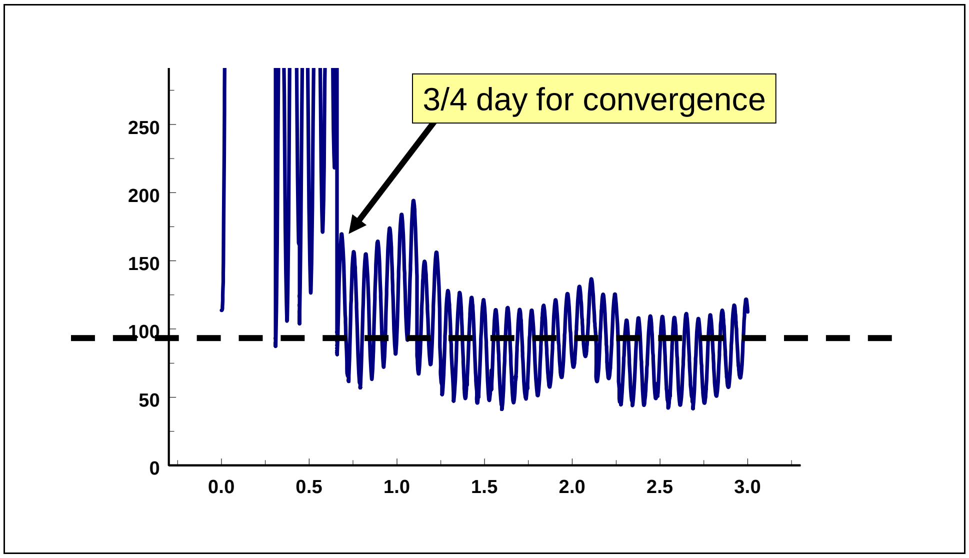

Fig 8: RSS of Position Sigmas for “Rule of Thumb” Tasking

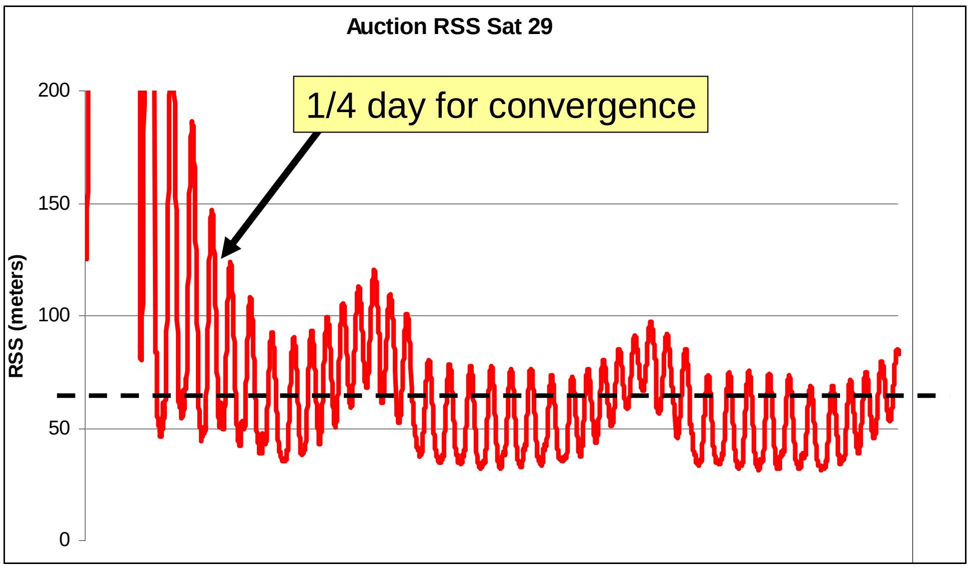

Fig 9: RSS of Position Sigmas Auction Method Satellite 29

Similar graphs can be obtained for each of the satellites and each of the Tasking algorithms. The challenge is to summarize the performance of each Tasking algorithm across all 296 satellites and to compare these algorithms to each other. The approach we have chosen is to find the median orbit error in an RSS sense for each of the Tasking algorithms and to provide the max and min values for each tracker, all on a five-minute grid. The resulting graph is a little busy (see Fig 10.).

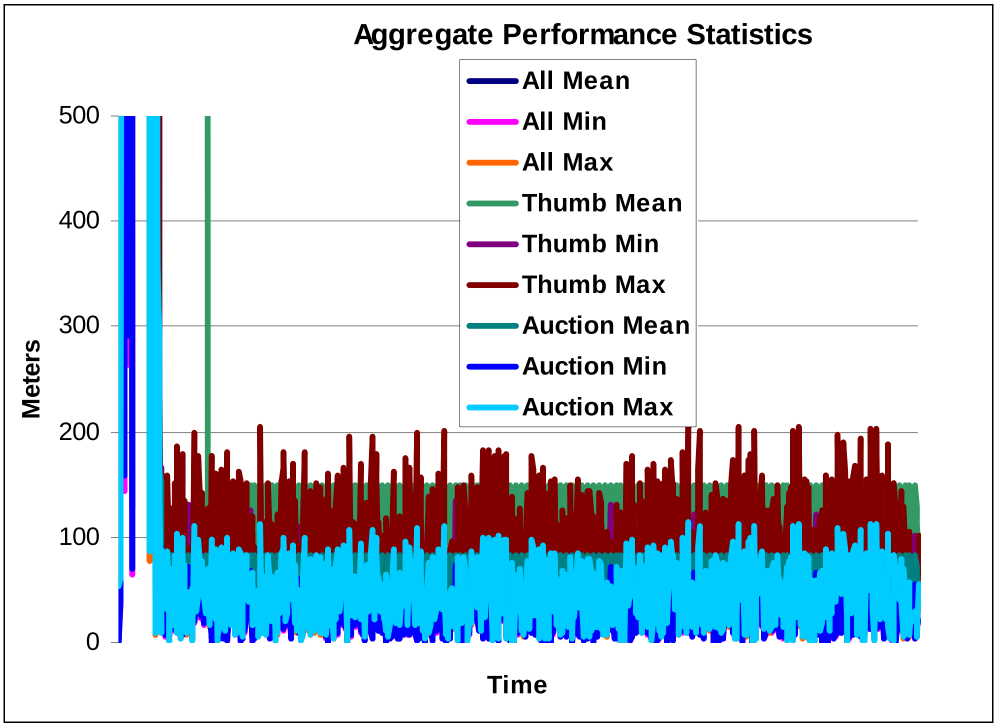

Fig 10: Summary Performance All Three Trackers

Altogether this auction, although somewhat naïve, illustrates the overall capability of the algorithm to improve Tasking over an intuitive method and it approaches the optimal capability defined by all possible tracking data.

7. FUTURE EFFORTS

Having had some success with finding bidding strategies for two diverse LEO agents, the next step is to devise bidding strategies for many other agents, for example rapidly decaying satellites, calibration satellites, rendezvous and docking scenarios, multiple satellites in a cluster (e.g. SES-ASTRA clusters in GEO), etc. It is also necessary to extend from LEO to HEO and GEO and to include optical sensors. Optical sensors will be challenging in that they are limited to nighttime-only operations, and may have local weather conditions to consider. If there is a meteorological report from the optical sites to the Central Tasker then the Auction should be able to take that into account and shift tracking responsibilities to the sites with good "seeing", and thereby dynamically maximize the coverage of the deep space satellites.

The next more complicated Auction is the Combinatorial Auction [2], where agents will bid on "packages" of resources (in this case a combinations of tracks). We will probably not attempt to develop a Combinatorial Auction algorithm, since the potential improvement in overall system operations over a simple discrete auction may not be worth the effort.

This effort has demonstrated that there is merit in modern optimization methods as applied to classical problems like Sensor Tasking. AGI has been approached by four different business partners, each with different optimization algorithm expertise, hoping to team to create a real advance in the state-of-the-art in resource management, using the SSN as an example. One of these vendors is Charles River Associates, who have considerable experience with Auction Methods. While it is not clear which of the four business partners AGI will team with (we may work with all of them) to pursue these algorithms, it is clear that there is a good reason that we will consider the Auction Method favorably.

8. ACKNOWLEDGEMENT

My first introduction to Auction Methods as a method for optimization in complex problems was during a conversation in 2007 with Dr. David Parkes [4], Professor of Mathematics at Harvard, where he told of sophisticated Auction algorithms that can keep track of the results of previous decisions and actually be used to detect agents that overstate their needs during the auction (e.g. they can "lie"), and can modify the bids based on consistent patterns of "lying". To me this was an enlightening moment, illustrating that Auction methods could actually be used in real-time operational settings to assign assets in a tactical situation and to prevent assets from being misplaced by overly "greedy" agents. The application of Auction Methods in military operations and in large scale civil disasters (Haiti, Katrina, Pakistan, etc) could have a tremendous impact on efficiency and responsiveness and are worth further exploration. I would encourage military and civilian agencies with complex decision making problems to consider funding the research being done at Harvard and other academic institutions to advance these techniques for real-world applications.

9. REFERENCES

1. Shoham, Y. and Leyton-Brown, K., *Multi-Agent Systems*, Cambridge University Press, New York, NY, 2009.
2 Crampton, P., et al, *Combinatorial Auctions,* The MIT Press, 2005.
3 Schwind, M., *Dynamic Pricing and Automated Resource Allocation for Complex Information Services,* Springer-Verlag, Berlin, 2007
4 Parkes, David, private conversation, 2007.