Globally Optimized Scheduling for Space Object Tracking

Richard Stottler[1] and Ryan Thompson[2]
Stottler Henke Associates, Inc., San Mateo, CA, 94404

Tracking of objects in earth orbit is an extremely important task for maintaining the safety and viability of manned and unmanned spacecraft. The sensors used to track these objects are mechanical and phased array radar and optical telescopes. Unfortunately, there are not enough resources to easily make enough observations to track every object's orbit to enough accuracy. This problem is exacerbated by the current space object tracking sensor scheduling process. A scheduler at each sensor site determines which opportunities and when within each opportunity to make observations (i.e. it schedules), without reference to what other sensors are observing the same object (and when) and which orbit metrics will contain the most and least errors. The schedule is not globally optimized, since scheduling occurs only locally without reference to what is assigned or scheduled at other sensor sites. By considering the effects of other observations of the same object by other sensors, specific sensors at specific times can be chosen that provide complementary observations of the object, which will reduce the object's orbit metric error covariance. We developed a scheduling algorithm that takes as input the space catalog and the associated covariance matrices and produces a globally optimized schedule for each sensor site as to what objects to observe *and when*. This algorithm is able to schedule observations that are more complementary, in terms of the precision with which each orbit metric is known, in order to produce a satellite observation schedule that, when executed, minimizes the covariances across the entire space object catalog. The algorithm has been prototyped and tested with the deep space objects in the space catalog, and the resulting orbit metric covariances calculated. Significant reductions in covariances were observed and are presented here.

Nomenclature

DS	=	Deep Space
GEO	=	Geosynchronous
LOS	=	Line of Sight
NP	=	Non Polynomial
TLE	=	Two Line Elements

I. Introduction

TRACKING of objects in earth orbit is an extremely important task for maintaining the safety and viability of manned and unmanned spacecraft. Space Track (http://www.space-track.org) provides two line elements (TLEs) for the approximately 20,000 manmade objects tracked in earth orbit that comprise the Space Catalog. The sensors used to track these objects are mechanical and phased array radar and optical telescopes. Unfortunately, insufficient resources exist to easily make enough observations to track every object's orbit to desired accuracy. This problem is exacerbated by the current space object tracking sensor-scheduling process. This process takes as input the current Space Catalog. Based on the characteristics of the object, the specifics of each sensor, and the geometry of each object's orbit with respect to each sensor site, a maximum probability of detection is calculated. Based on each object's priority, number of observations needed, and the maximum detection probabilities, each sensor is tasked to observe each object a specific number of times. A scheduler at each sensor site determines which opportunities and when within each opportunity to make observations (i.e. it schedules), without reference to what other sensors are observing the same object (and when) and which orbit metrics will contain the most and least

[1] President, 951 Mariners Island Blvd, Suite 360, San Mateo, CA 94404, AIAA Member
[2] Technical Staff, 951 Mariners Island Blvd, Suite 360, San Mateo, CA 94404, AIAA Contributor

errors. It does not necessarily schedule the same time or even the same opportunity that corresponds with the maximum probability of detection. There are issues with the current process:

1. The schedule is not globally optimized since scheduling occurs only locally without reference to what is assigned or scheduled at other sensor sites.
2. A globally optimized schedule would typically have more observations and a better number of observations for each object, with the same resources.
3. By considering the effects of other observations of the same object by other sensors, specific sensors at specific times can be chosen that provide complementary observations of the object, which will reduce the object's orbit metric error covariance matrix.

A scheduling algorithm was developed that takes as input the space catalog and the associated covariance matrices and produces a globally optimized schedule for each sensor site as to what objects to observe *and when*. This algorithm is able to schedule a better number of observations of each object with the same sensor resources and have those observations be more complementary, in terms of the precision with which each orbit metric is known, to produce a satellite observation schedule that, when executed, minimizes the covariances across the entire space object catalog. The results would be increased accuracy of the space catalog with fewer lost objects with the same set of sensor resources. The algorithm has been prototyped and tested in simulation and the resulting orbit metric covariances calculated. These results are presented in this paper.

II. Deep Space Objects

The Space Catalog is generally divided between Deep Space (DS) and Low Earth Orbit objects with the dividing line being a period of 225 minutes. Due to their greater distance, DS objects require sensors of greater sensitivity and precision, and longer tracking time, and there are fewer of these sensors, hence the heightened need for highly efficient and complementary scheduling. One especially important part of the DS region is the Geosynchronous (GEO) belt. (GEO can also refer to "Geosynchronous Earth Orbit"). GEO objects orbit once per day, meaning they are essentially fixed in longitude. This means that they have LOS to a large, relatively unchanging fraction of the earth all the time, but at a long range. Non-GEO DS objects tend to have long time windows of LOS to specific locations and the set of those locations changes relatively slowly over time. Sensor requests typically inherit the constraints for the types of orbiting object to which they relate. Resource management involves optimally tasking specific resources to meet the requests while obeying all of the constraints.

III. Scheduling Background

Scheduling problems involving meaningful resource selection are NP-Complete, meaning that the computational time of any algorithm guaranteed to produce an optimal solution is exponential with the number of those decisions. It is straightforward to build bad scheduling systems but difficult to build good ones. I.E. there are relatively simple algorithms, such as priority-based schemes, that will generate correct schedules but not very good ones in the sense that many requests will be rejected that could have been scheduled with a better algorithm. It is relatively difficult to build good scheduling systems because the NP-Completeness necessitates the use of good heuristics.

Many scheduling systems use a naïve approach to assigning resources to requests wherein the requests are fulfilled in priority order. If the current request cannot be scheduled, earlier-processed (and therefore higher priority) requests will not be reconsidered in order to try to accommodate the current request. In particular, if the current primary request cannot be scheduled, the request will be left unscheduled, even if it had been physically possible to shuffle some of the higher-priority scheduled requests to accommodate the current one.

Although a priority-order allocation guarantees that higher priority requests will be fulfilled before lower-priority ones, it is not optimal (and can perform very badly) and underutilizes important resources. Consider the following extremely simple example, which is therefore easier to use to illustrate this point, where three requests for the same time-period have been received, Requests 1, 2, and 3, respectively, and where the priority is highest for Request 1 and lowest for Request 3. Two different sensors exist, A and B, that happen to be situated to track two of those objects each, based on the orbits. Sensor A has the capacity to track one object and Sensor B can track two during the time in question. The probability of detecting Object 1 is slightly higher by Sensor A but Sensor B is also acceptable. Request 2 requires Sensor B and Request 3 requires Sensor A. Request 1, with the highest priority is selected first by the priority-based scheduler and since its highest probability of detection sensor, Sensor A, is available, this is what is allocated to it. The priority-based scheduler then considers Request 2, with the second highest priority and selects Sensor B, since it is the only one able to track Object 2. When the scheduler gets to Request 3, it finds that only Sensor A could track Object 3, but it has no more capacity. Meanwhile, Sensor B has excess capacity. Clearly, this is a suboptimal solution, since it was possible, following an algorithm that considered resource contention, to allocate a sensor to every request.

But perhaps a scheduler could be devised that systematically tried every possible solution and selected the best, and therefore optimal, one. In the example above, the number of possible solutions is 2 choices for request 1 times 1 choice for request 2 times 1 choice for request 3 = only 2 possible solutions. However, consider a set of requests consisting of only 30 simple requests where there are on average 4 meaningfully distinct choices (e.g. different sensors or time windows) for each request. This means that there are 30 distinct decisions with 4 choices each so the number of solutions is 4 x 4 x 4.... x 4 = 4^{30} = over a million trillion possible solutions, which is clearly impractical to systematically search. Moreover, the space catalog results in many, many more track requests per day than this, tens of thousands. This is the essence of NP-Complete problems. They require good heuristics to solve them in a close to optimal manner and simple methods for ordering the decisions (e.g. ordering the requests) can produce very poor results.

IV. Sensor/Covariance Background

An orbiting object's state at any given instant can be uniquely defined by either 6 orbital parameters (Eccentricity, Semimajor Axis, Inclination, Longitude of ascending node, Argument of periapsis, and Mean anomaly at epoch) or equivalently as its 3-D position and velocity vectors (6 total parameters). Radar systems tend to provide accurate range and range rate measurements (2 parameters) but are much less accurate as to angle and angular rate measurements (4 parameters with relatively high error). Meanwhile optical sensors provide good angle (2 parameters) and angular rate (2 parameters) but no range or range rate (2 parameters missing). For each sensing opportunity, a 6 x 6 covariance matrix can be calculated based on the geometry of sensor and sensed object and the accuracies of the sensor in different dimensions; the covariance matrix describes the probable volume of space for the sensed object (or, equivalently, the variances and covariances of the errors in measurement). These errors in the orbit metrics can be best reduced by complementary measurements. The degree to which potential additional measurements are complementary can be determined by calculating the covariance matrices for those candidate measurements and comparing how complementary they are to the existing covariance matrix. In general, if the first measurement was optical, then for another optical measurement to be complementary would require it to either be from an instrument far enough away (and not oriented in line with the previous sensor and sensed object) to achieve a cross-fix to the desired accuracy or to be at a substantially different part of the orbit. Conversely, if the first observation is optical, then almost any radar measurement which is not close to orthogonal to the optical measurement will be complementary (orthogonal in the sense that the line from the optical sensor to the sensed object is orthogonal to the line from the radar sensor to the sensed object). In particular, an immediate radar observation collocated with the optical sensor (or at least not orthogonal to it) will greatly reduce the combined covariance. The same reasoning applies if the first observation is radar and the second is optical. Any non-orthogonal geometry will produce a complementary measurement and greatly increase the orbit metric precision.

A typical radar used for tracking deep space objects has about 0.02 degrees of error in angle and anywhere from 26 to 160 meters in range error. For a geosynchronous satellite about 36,000 km from the radar, the angular error translates into about 13 km of error for both the along-track and cross-track measurements. So, radar is about two orders of magnitude more accurate in range than in angle, for objects at this distance. Meanwhile, an optical telescope, while giving no range information, is accurate to about 0.004 degrees in angle, which translates to about 2.5 km of error in angle for geosynchronous objects. Obviously, the most precise measurement would involve both an optical and a radar measurement in relatively quick succession; however, the deep-space sensors are currently over-taxed so this is not possible. Also while an optical measurement might take around 1 minute, a deep space radar measurement will take considerably longer. (Mitigating this somewhat is the fact that radar can operate at any time of day while the optical sensors only operate at night.)

So given limited sensor resources and that the large majority of measurements are made with optical sensors[1] with infinite range error, how should measurements be scheduled to most reduce the error in an object's position? First, consider what is currently done. Typically, objects are scheduled to maximize their probability of detection. The time usually occurs when the object is opposite the Sun from the sensor's perspective, essentially local midnight. Obviously one optical measurement like this gives no information about the object's orbit's size (semi-major axis) or eccentricity. The semi-major axis is related to the period of the orbit. So if another measurement is taken on a later night at the same place in the orbit, the period of the orbit and thus the semi-major axis can be determined. But almost no information about the eccentricity can be derived. Similarly, although the angular error is fairly low, it is not zero. Minor errors in angular velocity will mean the angle of the orbit with respect to the measurement will not be precisely known. If optical measurements are taken at the same point in the object's orbit, the object's location is reasonably precisely known at that point and exactly opposite that point but not known with much precision 90 degrees from that point (where both the angle and eccentricity errors will be greatest). Both of these issues point to taking complementary measures, optimally 90 degrees from the last measurement.

Measuring an object's location with real sensors, which have real error, effectively implies determining a volume where the object is likely to be. This volume around the measured point is described by the covariance matrix (a 6 x 6 matrix either in the position-velocity space or in the orbital parameter space). Far greater importance is placed on the position variances (3 of the diagonal elements) than on other parts of the covariance matrix. One measurement's covariance matrix can be propagated to the time of another measurement so that the two measurements can be combined. The result is approximately the intersection of the two volumes.

Consider two optical measurements taken relatively closely in time from the same or nearby sensors. Each spatial error volume will be long and thin and directed radially away from the sensor (since it is an optical measurement). Propagating the older measurement to the time of the newer one will result in a new combined covariance (volume intersection) not much smaller than either of the individual covariances because the angle between the two volumes is small (since they are taken at similar angles). The reduction in error is therefore relatively small. But consider two measurements taken at 90 degrees apart in the orbit. After propagating the first measurement 90 degrees along the orbit to the second measurement, the thin volumes are much more orthogonal than the previous case and the intersecting volume of two thin volumes oriented approximately orthogonal to each other is a much smaller volume (and therefore a much smaller covariance) than either of the two individual volumes.

Taking successive optical measurements closer to 90 degrees apart and farther away from 0 or 180 degrees apart will lead to smaller intersecting volumes and therefore smaller combined covariances. The measurements do not need to be precisely 90 degrees apart in the orbit. 70% of the benefit is realized at 45 degrees.

V. Scheduling Problem Setup

3160 deep space and geosynchronous satellites were considered for the purpose of this study. A NORAD element set from Apr 13, 2006 was the basis for the satellites' orbital parameters. Nine sensors were scheduled for Deep Space and Geosynchronous object tasking: 5 optical sensors and 4 radars. Sensor data was obtained from *Fundamentals of Astrodynamics and Applications* (Vallado, 2001)[2].

In order to model the Earth orbits of every satellite in our catalog, we used NORAD's SGP4 propagation code (available from www.celestrak.com). Updated by Lt Col. David Vallado, then revised again in 2008 (Vallado and Crawford, AIAA-2008-6770) [3], this is the most up-to-date version of publicly-available propagation code that most closely resembles that which is actually used by NORAD to generate Two-Line Elements. The code was tested and matched with example orbit propagations from papers in 1980 and 2008. Since we do not have the capability to take actual observations and measurements, for the purposes of this simulation we assume that the orbits generated by SGP4 are 100% accurate and represent a true measurement.

Ideally, a sensor track at a given time would consist of several individual observations, which would then undergo differential correction to correct the TLE to reflect the new information provided by the track. For our purposes, since we were not actually updating TLEs, we just needed to calculate the covariance matrix for each measurement, which we did using particle techniques. An observation at time t consists of a range, azimuth, and elevation value for radar sensors, or right ascension and declination for optical sensors. The assumed true values of the observation, as calculated by the look angle calculations and SGP4, are perturbed by Gaussian error added in accordance with the published sensor errors from Fundamentals of Astrodynamics and Applications[2].

Once these error-added observations are calculated, the Herrick-Gibbs method of orbit determination is used to produce a set of position and velocity vectors at time t. The covariance of the sample is found through the standard covariance calculation for each element of position and velocity, producing a 6x6 matrix from our set of 6-D vectors:

$$\mathbf{Q_n} = \frac{1}{n}\sum_{i=1}^{n}(x_i - \overline{x})(x_i - \overline{x})^T. \quad (1)$$

A program was written that generates visibility windows and look angles for a given satellite. First, the satellite's assumed position is calculated over the course of the time window specified, at 1 minute time steps. This produces a long array of position vectors generated with SGP4; all in Earth-Centered Inertial (ECI) coordinates. Then, for each sensor at each time interval, the sensor's position is calculated in ECI coordinates from its latitude, longitude, altitude, and the time. Next, the satellite's position is transformed to topocentric-horizon coordinates, with the sensor as the origin (see http://celestrak.com/columns/v02n02/). From here, range is $\sqrt{[r_S^2 + r_E^2 + r_Z^2]}$, elevation is $\sin^{-1}(r_Z / range)$, and azimuth is $\tan^{-1}(-r_E / r_S)$.

VI. Prototype Scheduler

We have initially concentrated on the covariance reduction issue, since this is particular to the space object tracking scheduling domain, and left other aspects of the scheduler unsophisticated. The prototype scheduler explicitly attempted to achieve approximately orthogonal measurements as its primary goal so it could be determined how much improvement in the covariances was possible.

The Scheduler uses the visibility windows to compute a day's schedule. Two main considerations currently exist that determine how satellites are scheduled. First, the prototype tries to schedule each satellite once per day (following the lead of reference 1), giving preference to satellites that have not been scheduled recently. Second, it attempts to capture different parts of the satellite's orbit each time a track is scheduled. It generates preferred optical scheduling times for satellites, each based on the last successful observation. The Scheduler is looking for a different part of the orbit: 90 or 270 degrees away from the last observed mean anomaly, an orbital parameter that corresponds to the satellite's position in its orbit. Due to a lack of optical resources relative to the demand, once they fill up, radar is used as an alternative.

Prototype Scheduler Algorithm:

* 24 hours is scheduled at a time, based on last successful sensor observation.
* Take in a set of visibility windows: satellite, sensor, time, duration, and last anomalies
* Goal: schedule different parts of the orbit while not starving any satellites
* The list of satellites to schedule is sorted by longest time since last observation
* Generate ideal optical window for each satellite
* Ideal optical window is at 90 or 270 degrees from last mean anomaly, in order to get observations at different parts of the orbit
* Attempt to schedule in order
* When there are conflicts, expand the window to 60 to 120 degrees or 240 to 300 degrees.
* If still are conflicted, try to schedule radar 24 hours from the last observation, +- 6 hours
* If still conflicted, try to schedule radar 24 hours from the last observation, +- 12 hours
* Once all satellites have attempted scheduling, write the schedule, go to the next day, read the schedule, repeat.

VII. Calculating Results

After each day, a schedule is written. The simulation was conducted over 3, 10, and 30 days with little difference noted in the results for 10 and 30 days, so most of the analysis was done with 10 days worth of data. The schedule file consists of a series of observation tuples: satellite, sensor, and time. We can generate ephemerides for each observation time, using the provided sensor errors in azimuth, elevation, range (radar) and right ascension, declination (optical) to provide a Gaussian distribution around the "true" values as determined by the SGP4 model. Turning these ephemerides into position and velocity vectors, via Herrick-Gibbs (FAA), we can generate a sample covariance matrix for each scheduled observation. We can then combine these observations with the covariance combination formula described by Fig. 1.

One specific issue was introduced by our simulation methodology, which treats the SGP4 model as 100% accurate for propagating orbits infinitely into the future. In reality, there are unmodeled effects and perturbations on the model and, for live satellites, there is potential for maneuvers to occur. Both of these factors increase with the length of time since the last measurement. So whenever a measurement covariance was propagated into the future (typically to combine it with another one), error proportional to the amount of time was added to the covariance matrix being propagated. Failure to add this error results in unrealistically small variances at the end of the simulated period, because every measurement only reduces the previously reduced covariance matrix. We chose to add an amount of error per day that was approximately equivalent to a pair of good cross-fixed optical measurements. This was done by adjustments in variance of 1.5 km^2 in position and .002 (km/s)^2 in velocity along the diagonal per day. So when a new measurement was made, the old covariance matrix was propagated to the time of the new measurement, and the variance described above was added and then combined with the new covariance matrix according to the following formula:

C1: Adjusted Old Covariance
C2: New Measurement Covariance
CF: New Combined Covariance (i.e. fused covariance)

Cf = C1 * Inverse(C1 + C2) * C2
Chaining these covariances together, we reach a final 3, 10, or 30-day covariance measurement for each satellite.

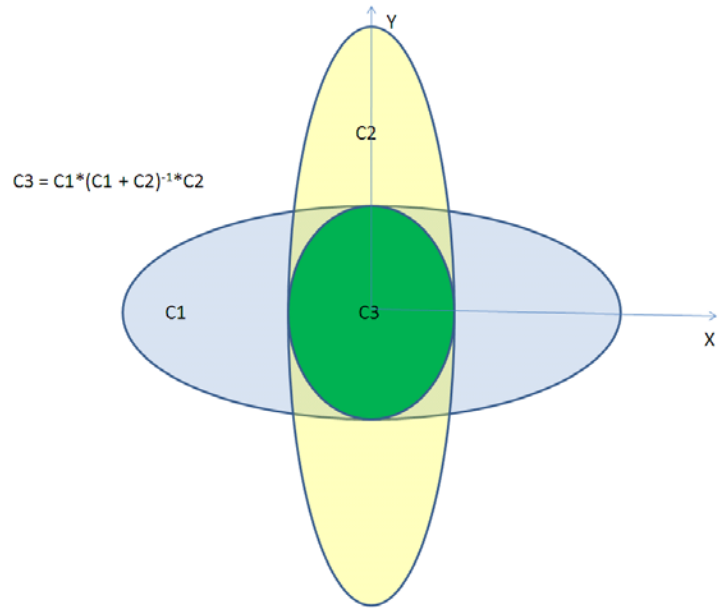

Figure 1. Covariance combination formula and corresponding sketch.

VIII. Results

3160 deep space objects from the catalog were tasked and scheduled by the prototype scheduler and compared to a simulation of the current method of scheduling for a 10-day period and the resulting covariances calculated in the position-velocity space at the end of the 10-day period. From these 6 x 6 matrices, the 6 variances were extracted and compared as described by Table 1 below.

Table 1. 3160 Deep Space Objects Final 10 Day Position/Velocity Variances.

	Prototype Average	Current Method Average	Current/ Prototype	Prototype Std Deviation	Current Std Dev.	Current/ Prototype
X $(km)^2$	3.36	14.43	4.3	3.26	45.27	13.9
Y	3.66	14.39	3.9	8.04	45.32	5.6
Z	3.11	8.08	2.59	4.46	21.32	4.8
Vx $(km/s)^2$	0.00148	0.00577	3.88	0.00130	0.00662	5.1
Vy	0.00151	0.00577	3.82	0.00131	0.00662	5.1
Vz	0.00141	0.00566	4.01	0.00130	0.00662	5.1

The first three rows refer to the x, y, and z components of the position variances in kilometer- squared units and the last three rows refer to the x, y, and z components of the velocity variances in kilometer per second squared units. The first row is the average of the variances across all 3160 objects for the 10-day schedule of observations output by the prototype scheduler and the second column is the corresponding average for the schedule produced by the current scheduling method. The third column is the ratio of the second column divided by the first and is a measure of how much better the prototype-produced variances were. The last three columns are similar to the first three but refer to the standard deviations of the variances instead of the average.

Clearly, having the scheduler purposefully schedule complementary measurements does significantly reduce the resulting uncertainty of the objects' locations after a ten-day period, producing variances 3 to 4 times better. The ratio of the standard deviations of the variances is even greater, indicating that the current method variances have a far greater spread, which is very undesirable. These high standard deviations for the current method indicate that there were some very bad measurement errors for some objects. One hypothesis to explain this is that the geosynchronous objects, although a very important part of the deep space regime, might have significantly higher measurement variances than the average because current method would tend to sense these objects in the same parts of their orbits each time. To investigate this hypothesis we ran the same statistics on the 674 GEO objects, resulting in Table 2 below.

Table 2. 674 GEO Final 10 Day Position/Velocity Variances.

	Prototype Average	Current Method Average	Current/ Prototype	Prototype Std Deviation	Current Std Dev.	Current/ Prototype
X $(km)^2$	4.95	29.73	6.0	2.17	72.60	33.5
Y	5.22	27.82	5.3	4.31	67.74	15.7
Z	4.13	9.17	2.22	5.78	16.11	2.8
Vx $(km/s)^2$	0.00191	0.00807	4.23	0.00228	0.00732	3.2
Vy	0.00193	0.00806	4.17	0.00229	0.00735	3.2
Vz	0.00168	0.00783	4.66	0.00231	0.00734	3.2

As hypothesized, the GEO objects had significantly higher average position variances for the current scheduler algorithms, especially in the X and Y dimensions. (The Z dimension is oriented "up" and the GEO objects tend to be near the equator, so large range optical errors from earthbound sensors tend to have a relatively small error in the Z-direction.) Some growth is expected, since angular measurements increase linearly with distance and GEO objects are significantly farther away than the average DS object. This can be seen by the slightly higher mean variances for the prototype. However, the X and Y components of the position variances for the current method double so that the ratio with the prototype is in the 5 to 6 range (or 4 to 5 if the Z component is included.) Even more significantly, the standard deviation of the current method variances is very high, indicating that there are some very large measurement errors, even after 10 days of observations. The prototype's standard deviations are 17 times better for the GEO objects. The GEO belt is an especially important deep space region and covariance- oriented scheduling greatly improves the accuracy of objects' positions and significantly reduces the variability of that accuracy (i.e. greatly reduces the worst-case situations).

These deep space results will not carry over to near-earth objects, where errors are reduced considerably due to their closer proximity and short orbit periods, which tend to give more variability in the position of the orbit where measurements are taken and the different angles between the sensor and object tend to give complementary angles.

IX. Future Work

The main purpose of the prototype was to verify that significant covariance improvement was possible and to somewhat quantify that improvement. However, the scheduling algorithm used was fairly simple in many ways. We are beginning work on a full-scale system and one of the main differences will be use of explicit covariance calculations and covariance fusion calculations as part of the scheduling decision process. Specifically, when selecting between competing observation choices, the choice that reduces the existing covariance the most should tend to be chosen, at least for the deep space objects where covariance is considered very important. Near earth objects may have other metrics considered important such as the number of observations per day or minimizing the time between successive observations. Especially important objects, which may need to be revisited frequently, may also need to focus on optimizing metrics different than covariance. Specific sensors, tailored to specific missions such as small debris sensing, will need to be freed from regular object tracking as much as possible to allow more time to the tasks that only they can perform. Finally, more sophisticated scheduling algorithms will allow more observations, and a more important (from each object's perspective) and better distribution of observations to be taken with the same resources.

X. Conclusion

Space object tracking is an important function for maintaining the safety of manned and unmanned spacecraft. The current method of scheduling observations of space objects is an enormous improvement over the way it had been done previously but is hampered by the fact that many of the existing sensors must perform their own scheduling without regard to what other observations other sensors are performing. This prevents the observation schedule from being globally optimized. However, the results of this paper show that infrastructure changes to allow for globally optimizing the schedule would have significant benefit in terms of greatly reducing the error in position with which a deep space object is known. This reduction is especially great in the very important geosynchronous belt.

References

[1]Miller, J., "Covariance Analysis for Deep-Space Satellites with Radar and Optical Tracking Data,", *Proceedings of the 2005 AAS/AIAA Astrodynamics Specialists Conference*, AIAA, San Diego, CA, 2005.

[2] Vallado, D., *Fundamentals of Astrodynamics and Applications*, 2nd ed., Microcosm, Inc., Hawthorne, CA, 2001.

[3] Vallado, D., Crawford, P., "SPG4 Orbit Determination," *Proceedings of the 2008 AIAA/AAS Astrodynamics Specialists Conference*, AIAA, San Diego, CA, 2008.