The All–Versus–All Low Earth Orbit Conjunction Problem[1]

Arthur Lue
MIT Lincoln Laboratory

ABSTRACT

A critical problem is emerging in the space environment. For the foreseeable future, we expect there to be a collision between a large object and a debris object every 400 days. But these collisions are hidden amongst 10,000s of close conjunctions per day and pose catastrophic threats to mission-critical payloads.

This report describes work addressing the problem of systematically identifying these needles in the proverbial haystack. Several key avenues are pursued: developing a computational scheme that can rapidly estimate conjunctions for large catalogs (∼100,000 objects); improving covariance analysis to effectively cull the number of encounters of critical interest; devising new covariance-based statistics such that, regardless of orbit-estimation quality, all true conjunctions of critical interest are successfully tagged; introducing an automated, adaptive tasking scheme to ascertain all potential collisions with a low false alarm rate; and providing quantification for tying conjunction-prediction performance to sensor-network performance and tasking parameters.

As a specific example, by taking a 50,000 object catalog in the most dangerous debris regime, we established a protocol where sensors take a series of observations over the course of 10 days to predict the 1.4 conjunctions of less than 10 m for the 11th day. Only 9 conjunctions per day need to be monitored to identify those 1.4 true conjunctions. The sensor-network load necessary is 3.8 times the loading necessary to maintain the latent day-to-day catalog.

1. INTRODUCTION

We no longer live on a world with a big sky. The Fengyun Chinese ASAT in 2007 and the unforeseen Iridium/COSMOS collision early in 2009 have brought the problem of all–versus–all conjunction into urgent focus. Space is a dangerous place: hundreds of thousands of objects as small as 1 cm in size *constantly* pose potentially catastrophic threats to mission-critical payloads.

One may perform a preliminary order-of-magnitude estimate to assess the scale of the collision risk of the existing space environment. Consider that there are a number of large objects of strategic concern to us. There are about 1000 such objects in the main debris zone (roughly 800 to 1000 km in altitude) and we estimate that their average cross-section is 10 m^2. The interactions of greatest frequency and of greatest concern are the conjunction of those large objects with the population of

[1] This work was sponsored by the Air Force under Air Force Contract FA8721-05-C-002. Opinions, interpretations, conclusions, and recommendations are those of the author and are not necessarily endorsed by the United States Government.

small, yet still dangerous, objects (whose size is > 1 cm). Using the best estimated numbers from the NASA–Haystack obseravtional debris studies, the expected collision rate is about 2.5×10^{-8} s^{-1}. This estimate implies that for the foreseeable future, one expects a collision between a large object and a > 1 cm objects *every 400 days*. Moreover, these collision will be hidden among *10,000s of* < 10 *km conjunctions per day*. Finding these catastrophic needles in the haystack will prove to be a daunting problem as we move into the future, this is the key problem we wish to tackle.

2. THE COMPUTATIONAL PROBLEM

At first blush, filtering seems the logical basis of any conjunction assessment: any given object almost never encounters any other specific object. Doing a brute force comparison of every potential collision pair would be a ridiculous proposition. Nevertheless, let us indulge the counterintuitive question, is filtering smart, effective or even useful? This question becomes rapidly less foolish particularly as the catalog gets increasingly larger. The fundamental issue is that filtering is an N^2–problem that still requires at least low-resolution propagation (where N is the catalog size). Even if that propagation is relatively cheap computationally, the cost still grows rapidly with N.

Let us reconsider the brute force approach: a straight propagation of the full catalog (an order N–problem), but with a time-step optimization. Given the catalog size, N, the total time period of regard, T, computation time per time step per catalog object, dt, and the propagation time step, $t_{\rm step}$, the total computation time for propagation is $\text{Time} = dt \cdot NT/t_{\rm step}$. One can miss conjunctions less than a specific distance threshold based on $t_{\rm step}$. The worst case: e.g., a 30 second time step corresponds to circumstance where a 100 km along-track distance of closest approach (DCA) can close to zero. There are about a million < 100 km conjunctions per day for a 10,000 object catalog. More generally, the number of conjunction refinements necessary is then given by $\text{Number} \sim 10^{-2} N^2 (T/\text{day})(t_{\rm step}/30\text{sec})^2$. For each conjunction, we need to perform a Newtonian search of the original $t_{\rm step}$ time interval to find the distance of closest approach. We expect to do about 10 time-step calculations, on average, for each conjunction of interest. In any event, this contribution is an N^2–computation whose coefficient needs to be estimated.

The total computation time is a sum of order–N and order–N^2 contributions. By optimizing the time step, one can reduce the overall N–dependence of this problem. This has the form:

$$\text{Time} = dt \cdot T \left[\frac{N}{t_{\rm step}} + \frac{c}{2} N^2 t_{\rm step}^2 \right] \ , \tag{1}$$

where c is the estimated coefficient of the conjunction refinement term. Optimizing the time step, we find $t_{\rm step}|_* = (cN)^{-1/3}$, suggesting

$$\text{Time}_* = dt \cdot \frac{3}{2} T c^{1/3} N^{4/3} \ . \tag{2}$$

What ostensibly is an N^2–problem using the traditional approach can be reduced to an $N^{4/3}$–problem through reconsidering the paradigm and optimizing the time step. Other items to note is that this calculation is trivially dependent on they type of model used for propagation. The overall computation time is simply proportional to the computation time per propagation time step.

We may refine this problem further. We previously took the number of refinement calculations to be dependent on the worst-case scenario where a given propagation time step, $t_{\rm step}$, led to some large number of conjunctions through a direct velocity calculation at a given distance. We know that the fastest velocity possible an orbiting body can take. Assume that both objects close on each other with that velocity. That sets the maximum threshold necessary. Given (for example) a 3 minute time step, a >2000 km maximum velocity threshold is necessary. One needs to refine an enormous number of potential conjunctions that get within 2000 km of each other.

Let us consider a more sophisticated alternative. We know that the objects almost never close with each other at the maximum possible velocity. One can establish a polynomial interpolation of the difference of the trajectories: $dx = dr_* + dv_*(t-t_*) + da_*\frac{1}{2}(t-t_*)^2 + \frac{1}{6}db_*(t-t_*)^3 + \cdots$, where t_* is the true time of closest conjunction. The governing equations for the evolution of the differential of the position is dominated by tidal forces, i.e., are proportional to dx itself $d^2(dx)/dt^2 = (GM/r^3)\left[dx - \frac{3}{2r^2}(dx \cdot x)x\right] \sim dx + \cdots$. In the cases of concern, $dx|_{t=t_*} = dx_*$ is parametrically small in natural units, in particular, much smaller than dt (also, $t_* \sim dt$) in natural units. Note that this analysis does not work in general, but only in cases where conjunctions of interest happen.

The natural dynamics imply that $dv_* \sim 1$ and $db_* \sim 1$ in natural units, but that dx_* and da_* are negligible. The dominant contributions to the polynomial interpolation are then

$$dx = dr_* + dv_*(t-t_*) + \frac{1}{6}db_*(t-t_*)^3 + \cdots \ , \tag{3}$$

where we are ignoring only $\mathcal{O}[(t-t_*)^4]$–terms. With these expressions, one can then use these expressions for a three-point interpolation and know that they are only in error to $\mathcal{O}(dt^4)$ in natural units. This represents a huge reduction in the threshold distance of necessary refinements, allowing a much smaller N^2–contribution to the computational problem.

Employing the techniques described in this section, one may efficiently enumerate all the conjunctions occurring for given catalog over a given period of time. Let us specify a particular scenario: SpaceTrack catalog, 2009 day 140 (May 20), full catalog analysis for entire day, identify exactly all the conjunctions within a given critical distance. We employ an SGP4 propagation for various reasons. First, SpaceTrack reports SGP4 element sets. Moreover, though actual results sensitively dependent on GP (analytic, low accuracy) vs. SP (numerical, high-accuracy) vs. truth, statistical results are robust against propagation details: statistical results consistent with a random spherical gas model of objects. For the full catalog 2009 day 140 (May 20): 3 conjunctions < 100 m, 274 conjunctions < 1 km, 27271 conjunctions < 10 km. For context, the average time between conjunctions per object is as follows: 5.9 years < 100 m, 24 days < 1 km, 6 hours < 10 km. These represent huge number even for the relatively small catalog size as it currently exists. We need to take steps to cull these conjunctions to focus on the ones of critical interest.

3. CULLING CONJUNCTIONS OF CRITICAL INTEREST

How do we proceed to cull the number of conjunctions of critical interest? First, we need to realize that distance of closest approach is not, at some level, a relevant measure of the danger of collision. We should be more concerned about where an object could be given our understanding of its state. Covariances are a measure of an object's orbit uncertainty, but care must be taken in order to use covariances are a reliable measure of orbit uncertainty. In this analysis, we use covariances based on the mean element sets based on insight gleaned from this prior work. However, an additional modification is employed where a canonical set of along-track variables (mean motion and mean longitude), as opposed to a non-canonical set (semimajor axis and mean longitude), form the root of the state-vector dynamics. This modification provides an additional improvement in achieving Gaussian containment over very long periods of time.

Association of two regions of uncertainty requires some amount of care. Intuitively, one would like to employ the logic of a simple normalized distance. Does the location of one object lie within k sigmas of the other object. However, the object in question itself does not have a well-defined location. It has its own uncertainty and its own set of k–sigma regions. How does one assess the *spatial* closeness of two ellipsoidal regions in state space? One needs to construct a more complex metric of association. Let us call this the *spatial joint k–distance*. This new distances provides a single number for conjunction alert for uncertainty. Consider the following. Find the minimum combined k–distance $k(x, v_1, v_2)^2 = k_1^2 + k_2^2$, where $k_i^2(x, v_i) = (\alpha(x, v_i) - \alpha_i^0)C_i^{-1}(\alpha(x, v_i) - \alpha_i^0)$, and α_i^0 is the estimated state of the i–th object, and C_i is its corresponding covariance. Note that the state α for object 1 versus object 2 is different. The spatial position is the same, but the velocities are different. The value k is then dependent on the common position x but on also on both velocities v_1 and v_2. The spatial joint k–distance value is when *all* three quantities are varied to minimize $k(x, v_1, v_2)$. The spatial joint k–distance is guaranteed to be positive and represents a normalized distance scale that characterizes whether two states with given uncertainties are likely to be at the same position.

Intuitively, the spatial joint k–distance characterizes the likelihood that two objects are in the same spatial location (but where the objects' respective velocities may, of course, be different). Schematically, the likelihood of collision goes like

$$P_{\text{collision}} \sim \int dx dv_1 dv_2 e^{-k(x,v_1,v_2)/2} \ , \tag{4}$$

where k is defined above. So long as k varies rapidly with the position and velocity parameters, one can estimate this probability using the method of steepest descent, and get the leading order behavior of the integral:

$$P_{\text{collision}} \sim e^{-k_{\text{joint}}/2} \ , \tag{5}$$

where k_{joint} is the minimum value of $k(x, v_1, v_2)$ subject to variations of x, v_1 and v_2. This minimum value is what we have called the spatial joint k–distance. The exponential factor computed is not the only contribution to the collision likelihood. There will be a multiplicative factor inversely proportional to a volumetric factor. However, that contribution is subdominant, and so the joint k–distance represents a good handle for characterizing the likelihood of collision; when $k_{\text{joint}} \gg 1$, it indicates very low likelihood of collision, whereas a small k suggests a conjunction of interest.

Consider a collision or a very close conjunction of interest. What would the predicted conjunction look like? How does the joint k–distance fare as a threshold criterion? Figure 1 shows the results. For the scenario that is essentially a collision, from 1000 Monte Carlo estimations,

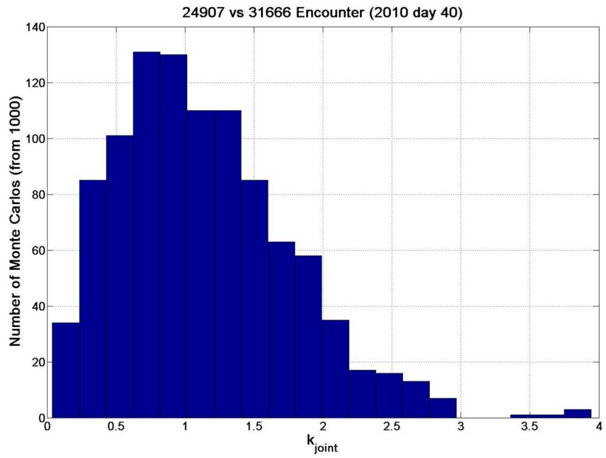

Figure 1. 24907 (IRIDIUM 22) versus 31666 (Fenyung debris). A scenario with a real "collision" or, in actuality, a 40 meter conjunction. The conjunction occurs at 2010 Day 40 at minute 89.629632. A three-day set of simulated observations is taken before the beginning of that day (with 1 km observation positional uncertainties). 1000 such Monte Carlo instances are run. The minimum value of the joint k–distance is compiled for all 1000 Monte Carlo runs.

conjunctions around the same time as the actual collision are predicted at various degrees of closeness. However, all of those conjunctions had $k < 4$ and almost all had $k < 3$; i.e., collisions will appear to be conjunctions with small joint k–distances. Alternatively, if the joint k–distance is greater than some specified value (say $k = 4$ or 5, or 10, etc.), then one can be assured with a specified degree of confidence that a collision will not happen. One can generalize this statement to a relationship between conjunctions of a specified distance and the joint k–distance threshold. There is a monotonic relationship between the two.

4. FULL–CATALOG SIMULATION

Can we use this k–distance statistic to positively identify the true conjunctions of critical interest using imperfect observations? The catalog used to test these conjunction ideas comes from the NASA 2030 Simulated Debris Catalog (objects $>$ 2 cm) with some modifications: 176,228 objects with 49,067 satisfying criteria apogee altitude $<$ 1200 km and perigee altitude $>$ 700 km and $0 < B^* < 0.1$. The following are the resulting truth statistics over the course of ten days: 141,572 conjunctions $<$ 1 km, 1411 conjunctions of $<$ 100 m, 18 conjunctions of $<$ 10 m.

We wish to see that, given simulated observations of these objects of a certain quality, can one predict these true conjunctions. For this simulation, we take the following criteria on sensor network performance: observation position uncertainty at $\sigma = 100$ m, three day fits with 11 sets of evenly spaced observations, ten day propagation into the future starting from the last observations. The simulation of observations, constructing estimated states, propagating those states and identifying the k–distances of the most critical predicted conjunctions (using the computational techniques

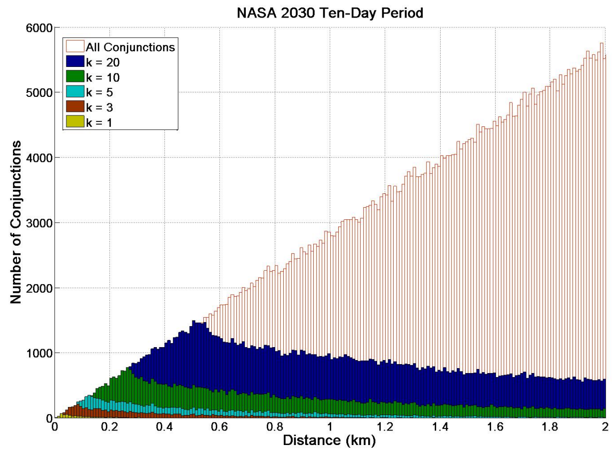

Figure 2. Culling of conjunctions using the joint k–distance.

prescribed in the last section) takes 7.5 days of computation using an 8–processor Intel Xeon CPU 5160 @ 3.00 GHz (4 GB RAM) machine using MATLAB code. State propagation was performed using SGP4 and original code for propagating accurate covariances.

The estimated statistics over the course of ten days are as follows: 141,716 conjunctions < 1 km, 1390 conjunctions of < 100 m, 12 conjunctions of < 10 m. How do these predictions line up with the truth? We use the k–distance to prioritize them. Figure 2 compiled the histogram of all the predicted conjunctions and their k–distances. Figure 3 shows the raw conjunction rate as a function of that k–distance; i.e., how many conjunctions per day occur with a k–distance less than a specified value? Note that this rate increases with time because the uncertainty volume increases with time.

There is a true conjunction out of every 10,000 that this algorithm simply misses. All of those result from an initial covariance which is pathologically large; i.e., they result from a bad catalog elset. These are easy to identify a priori, and presumably those orbits can be corrected by sensor network tasking. For initial covariances that are corrected from the initial orbit estimate, there is fewer than one failure for every 100,000 conjunctions (no failures are actually detected). This suggests that the algorithm for identifying and predicting conjunctions is sound, and now one only needs to set a k–distance criterion for conjunction alert.

We can see that for all conjunctions that lead to collisions (i.e., the true DCA is zero), we can take $k_{\text{joint}} < 2.5$ as our alert criterion. If we are concerned about conjunctions of < 100 m (i.e., conjunctions whose DCA is comparable to the observation uncertainty, $\sigma = 100$ m), then we need to take $k_{\text{joint}} < 6$ as our alert criterion. For the latter, there are 1411 such true conjunctions over the course of ten days, whereas the number of conjunctions that satisfy the k–distance alert criterion is over 27,000. What one can say is that taking that the $k_{\text{joint}} < 6$ criterion guarantees that in amongst those 27,000 conjunctions, all 1411 of the true conjunctions of concern are identified. However there is a large false alarm rate, roughly a rate that is ten times the true conjunction rate

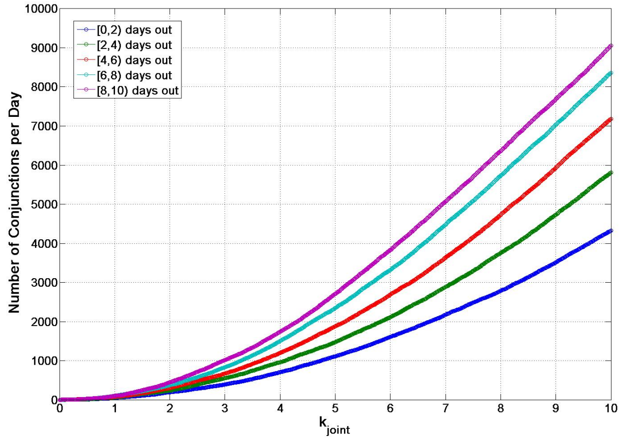

Figure 3. Conjunction rate of the NASA 2030 catalog as a function of the joint k–distance. Also note that for a fixed k_{joint}, the conjunction rate increases with time as the uncertainty volumes grow for each object. The conjunction rate grows approximately with the square of the joint k–distance.

initially, and a rate that grows to roughly 30 times the true conjunction rate, after ten days.

How do we reduce the false alarm rate? Simulations suggest the dependence is straightforward. The conjunction rate goes like $n = \gamma k^2 \sigma^2$, where $\gamma = 0.385 \times 10^4$ day^{-1}km^{-2} for the catalog and sensor system under consideration. To lower the conjunction rate, one can go to more sensitive observations, but we should do so in a way that accounts for the true distances of concern. By doing so, we are naturally led to the outline of an adaptive tasking scheme.

5. AN ADAPTIVE TASKING SCHEME

The k–distance statistic is very good for assessing whether a given conjunction is consistent with being a true collision, i.e., given orbit uncertainties, the two objects involved will occupy the same location, leading to catastrophe. The k–distance is an effective statistic for conjunctions whose DCAs are much smaller than the characteristic scale of the orbits' spatial uncertainty. However, we are often interested in circumstances when we are concerned about DCAs that are comparable to the orbit uncertainties. Particularly when we look to the automated tasking schemes. As the predicted epoch of conjunction gets closer and closer, we are increasingly concerned about improving our orbit uncertainties to the scale of a critical DCA. That critical DCA need not be zero distance. Indeed, because critical objects are often large, with cross-sections of many square meters, the critical DCA can by many meters as well.

When discussing the joint k–distance, we argued that it characterized the leading contribution to the collision likelihood. We may generalize this concept to conjunctions of a certain distance, δ.

The likelihood of two objects achieving a DCA less than δ is, again schematically,

$$P_\delta \sim \int dx_1 dx_2 dv_1 dv_2 e^{-k(x_1,x_2,v_1,v_2)/2} \left[1 - \theta((x_1 - x_2)^2 - \delta^2)\right] \ , \tag{6}$$

where θ is the step-function and where the exponent is given by

$$k^2(x_1, x_2, v_1, v_2) = (\alpha(x_1, v_1) - \alpha_1^0) C_1^{-1} (\alpha(x_1, v_1) - \alpha_1^0) + (\alpha(x_2, v_2) - \alpha_2^0) C_2^{-1} (\alpha(x_2, v_2) - \alpha_2^0) \ . \tag{7}$$

So, now the positions of each object is allowed to vary freely, subject to the constraint that they be within δ distance of each other. Again, using the method of steepest descent, the leading order behavior of the integral becomes $P_\delta \sim e^{-k_\delta/2}$, where k_δ is the minimum value of $k(x_1, x_2, v_1, v_2)$ subject to variation of all the parameters subject to the constraint that $(x_1 - x_2)^2 \leq \delta^2$. We call this new distance statistic k_δ the *extended spatial joint k–distance.* Again, when $k_\delta \gg 1$, it indicates a very low likelihood of the two objects having a DCA less than δ.

By compiling the raw conjunction rates as a function of the parameters δ, σ and k, we find that the conjunction rate goes as $n = \alpha\delta^2 + \beta\delta k\sigma + \gamma k^2\sigma^2$, where α, β and γ are set by the catalog and the sensor system. For our simulation (on the first day): $\alpha = 1.4 \times 10^4$ day^{-1}km^{-2}, $\beta = 1.95 \times 10^4$ day^{-1}km^{-2} and $\gamma = 0.385 \times 10^4$ day^{-1}km^{-2}. The first term is just the true alarm rate; i.e., when $\sigma \to 0$, for any fixed k alert criterion, the conjunction rate approaches the true conjunction rate. This implies no false alarms. This is the result one would expect as the observations become increasingly accurate. The third term is just the original joint k–distance rate, again what one would expect when $\delta = 0$.

Another way of representing this expression is

$$n = n_0 \left[1 + \beta_1 \frac{k\sigma}{\delta} + \beta_2 \frac{k^2\sigma^2}{\delta^2}\right] \ , \tag{8}$$

where $n_0 = \alpha\delta^2$ is the true alarm rate, and the term in brackets is the false alarm ratio, where $\beta_1 = 1.4$ and $\beta_2 = 0.275$ for our particular simulation.[1] Figure 4 depicts this function.

This relationship determines the false alarm rate as a function of the observation accuracy, our alert criterion value, k, and our distance of critical interest. This allows us to now identify how many objects we can be expect to track as a function of the amount of sensor resources (as characterized by the observation accuracy) we are willing to put into those objects. Let us flesh out that relationship.

The table depicted in Figure 5 highlights an example tasking sequence that can adaptively and automatically monitor the most critical conjunctions that occur every day. The scheme relies on taking several distinct sets of tracks on a given day that are relevant to a series of days in the future. For example, take the high-density ∼50,000 object catalog that was used to analyze the future NASA 2030 environment in the high-density regime. Every day, a latent number of observations are taken for *every* object in the catalog. In order to minimize the loading on the SSN, we take that latent observation accuracy to be something quite low.

[1] We know that β_2 has a time-dependence of roughly $\beta_2(t) = 0.275(1 + t[\text{days}]/5)$. By dimensional analysis, one would then expect a time dependence for β_1 to be such that $\beta_1(t) = 1.4(1 + t[\text{days}]/5)^{1/2}$.

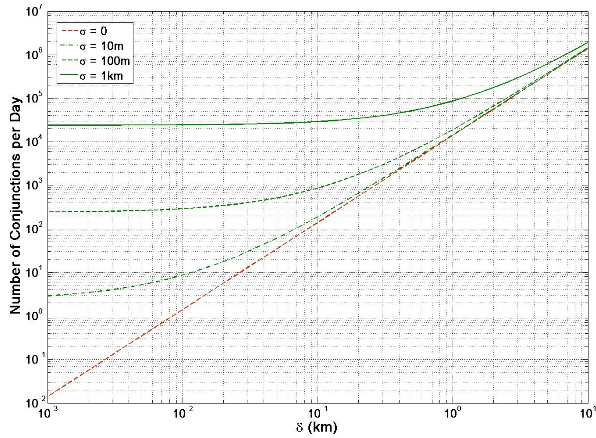

Figure 4. New conjunction rate statistics. For differing observational uncertainties, σ, the conjunction rate varies with the critical distance, δ for a fixed alert criterion $k_\delta < 2.5$. The dashed red line represents the true conjunction rate if δ would represent d_{truth}. Note that as $\delta/\sigma \to \infty$, the conjunction rate approaches the true rate. The difference between the green curves and the red curve is the false alarm rate: the number of extra conjunctions that, given orbit uncertainties, are indistinguishable from the conjunctions of concern.

Target Obs Accuracy	Number of Objects Tracked (with evol)	Degrees of Track per object
320m (latent)	50,000 (all)	.25
225m	15,000	.25
160m	7200	.5
113m	3500	1
80m	1650	2
57m	800	4
40m	390	8
28m	190	16
20m	98	32
14m	53	64
10m	29	128
7m	?	?
Obs Equivalent Sum Total	185,000	

Figure 5. Automated tasking scheme. On any given day, each row indicates the number of objects needed to be tracked that day at a given obs accuracy. The total at the bottom represents the number of obs equivalent to the latent obs accuracy (320 m) needed to be taken on that day; i.e., to get an observation of 225 m accuracy, need to take two 320 m accuracy observations, to get an observation accuracy of 10 m, need to take 512 observations of 320 m accuracy, etc. The "Degrees of Track per object" column is the number of degrees of track needed to produce that level of accuracy assuming the tracking beams are operating in a simple manner where 100 m accuracy come from 1 degree of tracking and that doubling the number of degrees of track improves accuracy by $\sqrt{2}$.

In fact, we set that accuracy threshold to be such that when the states for those objects are projected to a target day that is 11 days into the future, roughly half of them (here, 15,000 objects if one includes the growth of the orbit uncertainties over that time) will come within the specified δ–threshold (with $\delta = 10$ m) of some other object on the 11th day. This means the next day, we need to improve the orbit accuracy of those 15,000 objects. We task tracks to be taken on those objects of the *latent* accuracy and combine those with the observations already that day taken on from the latent set of observations, to improve to orbit accuracy by a factor of $\sqrt{2}$. When those 15,000 objects are propagated out the target day (10 days in the future), now only 7200 objects will come within the new δ–threshold. On the subsequent day, a new set of tracks are taken on those 7200 objects with an accuracy that is a factor of $\sqrt{2}$ better than the latent accuracy. When those are combined with the prior *improved* accuracy from the previous day, we get orbits that are two-times more accurate than the latent orbit accuracy. These orbits are propagated out 9 days, and there only 3500 objects come within the new δ–threshold.

This sequence is carried out until the day before the the target day (11 days after the original day in question), where 29 objects need to be tracked with 10 m accuracy. What results are 17 objects in 8.7 conjunctions that now are consistent with being within < 10 m conjunctions the next day. Only 1.4 of those conjunctions (on average) will actually be < 10 m, but with the δ–threshold set as specified, over 99% of the time, those actual < 10 m conjunctions will occur within the predicted set. As an aside, to get this number 8.7 to be closer to 2 predicted conjunctions (i.e., predict 2 conjunctions with 1.4 being the actual < 10 m conjunctions), need to improve obs accuracy by a factor of 10. Even for only tracking 17 objects, this may be beyond the scope of the current SSN, given sensor constraints and bias issues (i.e., getting meter scale obs accuracy).

We have followed a specific set of conjunctions predicted on a given day, starting from 11 days beforehand. Conversely, on any given day, we need to take the observations of every set of tasked tracks for the predictions 11 days in the future, 10 days in the future, 9 days, etc. Each set is a different set of objects (in general), and combined they represent about 185,000 observations per day, observations equivalent to the latent observation quality. The total number of obs equivalent is to take the total number of degrees of tracking divided by the degrees of tracking needed for a single ob (here assumed to be 0.25 degrees). Tasking load requires taking about 3.8 times as many obs as is necessary for the latent catalog. However, the load is not evenly distributed: a small number of objects demand large fractions of that load.