

APPROACHES TO EVALUATING PROBABILITY OF COLLISION UNCERTAINTY

Matthew D. Hejduk,* Lauren C. Johnson†

While the two-dimensional probability of collision (P_c) calculation has served as the main input to conjunction analysis risk assessment for over a decade, it has done this mostly as a point estimate, with relatively little effort made to produce confidence intervals on the P_c value based on the uncertainties in the inputs. The present effort seeks to try to carry these uncertainties through the calculation in order to generate a probability density of P_c results rather than a single average value. Methods for assessing uncertainty in the primary and secondary objects' physical sizes and state estimate covariances, as well as a resampling approach to reveal the natural variability in the calculation, are presented; and an initial proposal for operationally-useful display and interpretation of these data for a particular conjunction is given.

INTRODUCTION

The development and operational integration of the probability of collision (P_c) calculation, which first entered the industry through the work of Foster and Estes¹ and has since been refined by a number of practitioners,² constituted a major event in conjunction assessment risk analysis. Before the P_c , risk assessment was based solely on the miss distance at the time of closest approach (TCA) between the two conjuncting objects, and this datum by itself is not particularly helpful because it does not take account of the expected estimation error in the state estimates for the primary and secondary objects—if both state estimates are very uncertain, then it is quite unlikely that the two objects will actually achieve so close a miss. The P_c calculation quickly replaced miss distance as the principal risk assessment parameter because it does explicitly consider the two objects' state estimation uncertainties and produces an actual risk parameter—that is, a probability.

As formulated, this calculation generates only a point estimate of a probability; it does not naturally produce confidence intervals or a probability density function (PDF) as part of the output, from which a user could ascertain the uncertainty in the calculated parameter. To be sure, uncertainty information is an input to the calculation in the form of the primary and secondary object covariance matrices; but there are meta-uncertainties about the covariance matrices (*i.e.*, uncertainty about the uncertainty), as well as uncertainties in determining the sizes of the two objects in formulating the hard-body radius (HBR) for the P_c computation. Additionally, there is a

* Chief Engineer, NASA Robotic CARA, Astrorum Consulting LLC, 10006 Willow Bend Drive, Woodway, TX 76712.

† Analysis Lead, NASA Robotic CARA, Omitron Inc., 555 E. Pikes Peak Ave, #205 Colorado Springs, CO 80903.

natural sampling variability in the calculation for which the nominal Pc calculation is the expected value but which has its own null distribution governed by the amount of determinacy in the two input covariances. One would like to be able to include all of these perturbations in the calculation to produce a probability density of Pc values. A probability density, if formulated and presented correctly, can more fully reveal a conjunction's sensitivity to different input parameters by showing the change in probability density as different input parameters are altered. It can also allow a more meaningful comparison of the result to a user's single-parameter threshold for taking remediative action. If, for example, the nominal value is below the remediation threshold but a good bit of the PDF lies above this threshold, a conservative user may opt for remediation even though the Pc point estimate would not have suggested it.

The present analysis effort has attempted to parameterize covariance and HBR uncertainty, and to develop a resampling technique to reveal the natural variability in the Pc calculation, in a manner that can be folded into an overall Monte Carlo approach to produce a single output probability density that will take cognizance of all of these factors. The purpose of the construct is thus to give users a sense of the dynamic range of the Pc and allow them to make risk assessment and remediation decisions based on percentile points from this probability density rather than the more opaque single-point expected value. Test data to date indicate that the median value for the PDF often lies significantly below the expected value (not a surprising result for a right-skewed distribution), so users will also be able to see how much the nominal Pc differs from the median value, the latter of which being a more natural expression of central tendency. What this proposed construct will not do, or if it does it will be more by serendipity than design, is to provide a banded prediction of the nominal Pc value that will be generated from the next Conjunction Data Message (CDM). It is to be hoped, of course, that much of the time it will in fact serve this role; but changes in satellite tracking sources and amounts, as well as unreliable atmospheric density predictions, can substantially alter the state estimates and covariances, and thus the Pc, in a very much non-predictable manner. In such cases the new CDM presents the user not so much with an update but essentially a new problem, which is why the trending of conjunction risk assessment parameters is a difficult business that must be very carefully and modestly pursued.

The following sections detail the parameterization of the covariance and HBR uncertainty, explain the resampling technique developed to reveal the natural variability in the Pc calculation, discuss the Monte Carlo approach used to unify these elements into a single results-generation session, and present a proposed display of results to a user.

COVARIANCE UNCERTAINTY

Of all of the inputs to the Pc computation, the covariances of the primary and secondary objects (and especially of the latter, as it is typically much less well determined) introduce the most uncertainty—a somewhat ironic statement perhaps, as it speaks to the uncertainty of an estimate of uncertainty. In certain situations, uncertainty in position covariance size can have a substantial effect on the calculated Pc. The sensitivity of Pc to covariance size is illustrated in Figure 1, which shows the relationship between Pc and the covariance-size-to-miss-distance ratio. For simplicity, the combined covariance is presumed here to be spherical; but the same phenomenon is observed when it is allowed to assume an ellipsoid in an arbitrary orientation.³

The x-axis gives the ratio of the 1-sigma combined covariance size to the conjunction miss distance, and the y-axis the base-ten logarithm of the ratio of the calculated Pc to the maximum value the Pc can assume over the entire x-axis range. As one moves from larger to smaller covariances (as determined by the ratio previously described), the Pc slowly but steadily increases until it reaches a maximum (at a ratio value of $1/\sqrt{2}$) and then very rapidly drops off after that. If one

is situated on the left side of the peak, covariance uncertainty can affect the P_c somewhat; but if one is situated on the right side of the peak, then even small changes in the covariance size can drive very large changes in the covariance-to-miss-distance ratio and thus the resultant P_c . Since in many situations it is not precisely clear where one is situated on the curve (given uncertainties in both covariance and future miss distance), it is important always to give proper attention to potential uncertainties in the covariance and if possible to account for them.

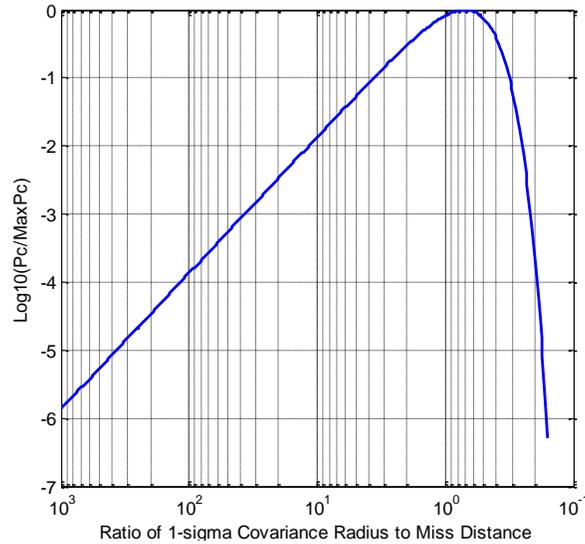


Figure 1. P_c Decrement as a Function of Ratio of Covariance Size to Miss Distance.

A straightforward and practical solution to this problem has been implemented in the JAC (JAVA for Assessment of Conjunctions) software package offered by the Conjunction Analysis and Evaluation Service: Alerts and Recommendations (CAESAR) program sponsored by the French Centre National d'Études Spatiales (CNES). A range on the “realism” of the primary and secondary covariances is determined, meaning the range of multiplicative factors by which the covariance might have to be scaled in order to represent the true error volume realistically. In the past, the range typically applied was from 0.2 to 5, although the ranges currently used are believed to be somewhat smaller. The primary and secondary covariances are, in nested fashion, scaled by factors in this range, the P_c under such conditions is computed, and the results are displayed in trade-space format, an example of which is given in Figure 2 below.⁴ In this figure, the scale factor for the primary is given on the x-axis and the scale factor for the secondary object on the y-axis. The unscaled case occurs where these two axes intersect, and the color of the square indicates the resultant P_c ; here, the brown color indicates a P_c of $\sim 5E-04$. One can see that increasing the size of either the primary or the secondary covariance by not a large amount at all (a scale factor of 2 or greater for either or both) puts the situation into the red zone, in which the P_c is greater than $1E-03$. If an owner/operator threshold for action were $1E-03$, then reasonable uncertainty in the covariance could allow the actual P_c to exceed $1E-03$ and thus be actionable even though the nominal value ($5E-04$) falls below this level. Considering this kind of uncertainty in the covariance is thus a prudent CA risk assessment activity.

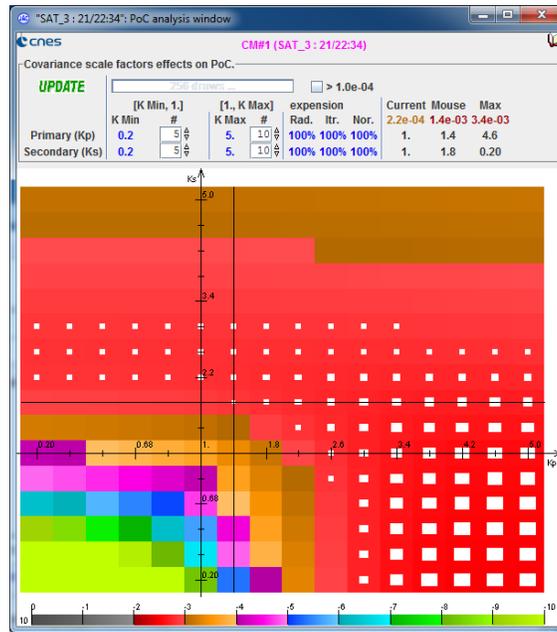


Figure 2. JAC CoPoC Sensitivity Display.

This particular approach does, however, encounter two types of limitations: it is not necessarily following a realistic spread of scale factors for the particular conjuncting satellites (0.2 to 5 may be far too large or perhaps even too small a span in a given situation), and it inherently presumes an equal likelihood that the real covariance could assume any of the sizes given in the span. One approach that could be somewhat helpful would be to solve for a scale factor (or full covariance matrix correction) that would, on average, transform the covariance to its proper size; this would at least debias the scaling procedure, allowing it to start from a properly-determined mean value rather than the arbitrary starting point of no corrective scaling. Approaches by Vallado and Seago,⁵ Cerven,⁶ and Hornwood *et al.*⁷ illustrate different methods for generating and certifying such corrections; unfortunately, these approaches as proposed do not give a probabilistically-enabled set of scale or transformation figures that can be used to assess the uncertainty in the covariance or, equivalently, the likelihood that the covariance could require a scale factor of a certain size*; but it is a simple and direct conceptual enhancement of any of these methods to produce output that can be used for this purpose.

This enterprise is simplified considerably by a utility that was recently implemented at the Joint Space Operations Center (JSpOC) to evaluate the covariance propriety of the Special Perturbations (SP) catalogue. Some years ago, functionality was introduced to build reference orbits for every catalogued object, much in the way precision reference orbits are constructed from satellite laser ranging data.⁸ The presence of these reference orbits enables a second utility to analyze every SP vector that is produced by propagating the vector and its covariance forward to propagation points of interest, comparing the propagated state to the reference state and thus generating position residuals, and combining the propagated covariance with the reference orbit co-

* Cerven (2013) does outline an approach to providing covariances sized at upper and lower percentile points of the expected full range of values, but this method is not particularly well suited to producing a large number of values across the whole distribution, which would be necessary for Monte Carlo applications.

variance to generate a combined covariance at the propagation point. Recent improvements to the process have instituted controls that greatly reduce the likelihood that the reference ephemeris and the propagated vector will share any observational data, as well as preserving the reference ephemeris covariance so that it can be combined with the propagated covariance to generate a total uncertainty of comparison. The availability of these products enables the straightforward calculation of covariance realism statistics.

The particular realism approach to use is here chosen in order to generate the particular final product for the present software tool, namely a set of scale factors that can probabilistically represent the covariance's uncertainty. If the products described in the above paragraph are generated for perhaps a year's worth of data on a satellite, one could have several hundred sets of position residuals and propagated combined covariances for a particular satellite at a particular propagation state. For each such pairing of position residuals (represented as a vector ε) and combined covariance C , one can compute the Mahalanobis distance

$$M = \sqrt{\varepsilon C^{-1} \varepsilon^T} \quad (1)$$

If the state estimate component errors are normally distributed (a necessary assumption if the covariance matrix as presently constructed is truly to represent the state estimate error), then the squares of the Mahalanobis distance calculations (M^2) should constitute a chi-squared distribution of three degrees of freedom.⁹ This principle can be used to transform the set of Mahalanobis distances into a set of scale factors that represents the covariance uncertainty. Each Mahalanobis distance value, it can be said, is itself the ideal scale factor for the covariance against which it is calculated, as it represents the factor by which the covariance must be pre- and post-multiplied in order to create the 3-DoF chi-squared expected value of 3. However, it is not correct, or is at least a distortion of the situation, to presume that, for a single evaluation of a particular covariance, the calculated Mahalanobis distance is a durable scale factor that could be applied to the entire family of covariances for that object at that propagation point to make all such covariances realistic. While the Mahalanobis test statistic does have an expected value, this is merely the hypothesized mean of an entire distribution of Mahalanobis test statistics; in such a situation one does not expect every sample to equal the expected value.

A simple approach can transform the Mahalanobis values into the desired uncertainty statistics. Presuming that a reasonably large sample of these Mahalanobis values for a particular object is available, one can rank-order the set and compare each calculated Mahalanobis value to the 3-DoF chi-squared cumulative distribution function value for the appropriate percentile point; the ratios of these two values do create scale factors that can be used to represent covariance uncertainty. As an example, suppose 100 M -values for an object are available; their squared values are calculated and then rank-ordered, and the 80th value has a value of 7.5. The 3-DoF chi-squared distribution at the 80th percentile has a value of 4.61. The ratio of these two values is 1.62; this is the value by which the covariance should be multiplied (or, if one prefers, pre- and post-multiplied by 1.27, which is the square root of 1.62) to be sized realistically, based on this one comparison in the context of the rest. The set of 100 such corrections can be used as a set of covariance realism scale factors from which to draw for Monte Carlo applications. Because they are a result of a large sample, these values already follow an appropriate distribution; so drawing from this sample set randomly will give proper probabilistic representation of the likelihood of particular scale factors. One can either perform these draws from the entire sample set or, if computation time becomes an issue, divide the sample into equiprobable bins and use a gridded combinatorial or Latin Hypercube approach.

HARD-BODY RADIUS

After the nominal miss distance and combined covariance, the hard-body radius (HBR), which governs the area in the conjunction plane over which the position probability density will be integrated, is the remaining input to the Pc calculation. So long as the combined covariance is reasonably larger than the nominal miss distance and this miss distance itself reasonably larger than the hard-body radius, the Pc follows a linear relationship in log-log space with the hard-body radius¹⁰; so one can perform the more computationally-burdensome parts of the 2-D Pc calculation (which can involve numerical integration) and then apply any HBR perturbations to that result set as a scaling parameter. This obviates the need to perturb the HBR parameter within a larger Monte Carlo construct, thus increasing computational efficiency. There are many approaches to determining the HBR for any given conjunction, each having its own advantages and disadvantages; the particular methods planned and implemented for the Pc Uncertainty framework will be surveyed and explicated here.

A common approach is to circumscribe the primary and secondary objects each with a sphere and define a HBR sphere that is the sum of these two radii; a diagram illustrating this approach is given in Figure 3 below.

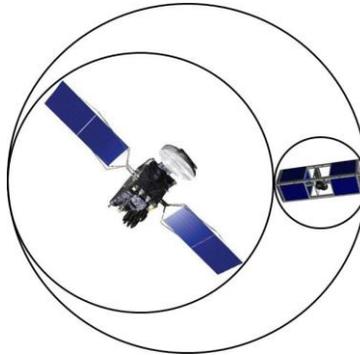


Figure 3. Primary, Secondary, and Combined Bounding Spheres.

Because the size of the secondary is generally not known, one usually makes a very conservative guess at its size, and then often adds even further conservatism to the final value by choosing an overall HBR that is larger than the sum of the postulated radii. If one elects to proceed in this manner, then one has specified a fixed HBR¹¹; and there can thus be no perturbation of this input within the overall calculation.

Primary Object

In trying to improve the estimate for the HBR of the primary, one can make a simplified geometric model of the spacecraft (or use a high-precision model, if it exists) and project the three-dimensional model into an arbitrarily-situated plane; if the full array of possible aspect rotations of the spacecraft are instantiated and projected into the plane and the projected area of each calculated, this set of projected areas forms a PDF of the possible different presented areas in the conjunction plane. One can then take percentile points from this distribution and construct in the conjunction plane a circle with the equivalent area to use for the 2-D Pc calculation; one can also take a step back from this and simply define a circumscribing circle for each projection. Figure 4 shows as an example the geometrical configuration and results set for the Orbiting Carbon Observatory 2 (OCO-2) satellite. The left pane shows the satellite in an arbitrary projection with a circumscribing circle drawn about it. The blue circles show the minimum, maximum, and average size of the circumscribing circles about each of the realized projections of the satellite into a

plane. Because the satellite (with solar panels considered) is long and relatively thin, a large range of circumscribing circles is possible: the flat projection shown in the left pane of the figure produces a large circumscribing circle, whereas a projection in which the long width of the satellite is normal to the projection plane will be very small and produce a small circumscribing circle; the ratio of smallest to largest circle is about 1 : 4.5, and the average size of circumscribing circle (across all of the possible projections into the plane) is 3.6, which is much closer to the maximum than the minimum value. The green circles represent the size projected areas themselves, here shown as areas of circles rather than the irregular shapes of the actual projections. They are substantially smaller, and their overall range of values is also tighter.

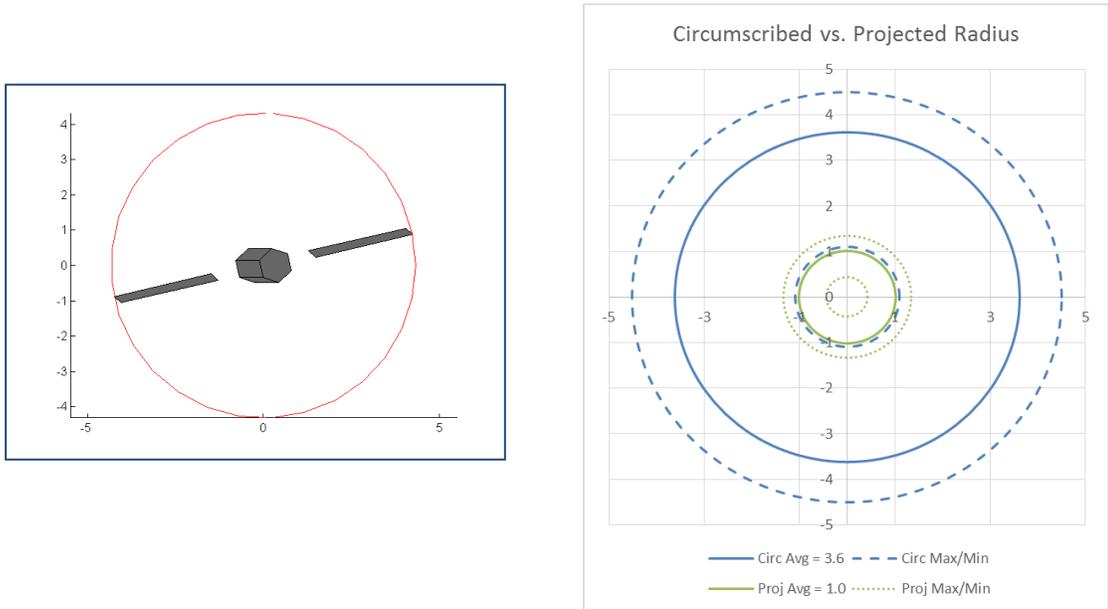


Figure 4. Projection Example and Projected Circumspection and Area Comparison for OCO-2 (sizes denominated in meters).

“Ball-and-stick” models that make use of properly-dimensioned simple shapes can perform quite well in comparison to higher-fidelity models. For example, a cylinder and flat-plate model of the Hubble Space Telescope (HST), constructed as a proof of concept, produced an average projected area that differed from that of a high-precision model value¹² by less than 5%. The amount of HBR reduction wrought by the use of such a technique can be considerable. In the case of HST, the circumscribing sphere has a radius of 8.1 m, but the radius of the circle that has an equivalent area to the median projected area of the spacecraft is 4.3 m. In any case, the full PDF of these projected areas can serve as a source of probabilistic perturbations to the HBR.

If the steering law of the spacecraft is known, then it is possible to determine the precise aspect orientation of the spacecraft at TCA. Since the conjunction plane is also known, if a geometric model of the spacecraft has been produced the model can be projected directly into the conjunction plane and thus the actual projected area for this particular event (for the primary, at least) can be determined precisely. One can convert this to a convenient shape to facilitate its combination with the size of the secondary object, or one can try to take direct cognizance of the particular shape by performing a contour integral about the object’s boundary as part of the 2-D Pc calculation. This latter approach must address the additional difficulty of how to grow the projected shape to account for the secondary object.

Secondary Object

Estimating the size for the secondary presents different challenges. If the secondary is an intact spacecraft, it is in principle possible to obtain dimensions for it and thus perform an average projected-area calculation following the methodology described above. Intact payloads are too numerous and varied to make determining individual models practical, so one might quickly assign one of a few basic “canned” models and use the projected-area PDF associated with that model type, after appropriately scaling for the payload dimensions. The same procedure can be followed with rocket bodies, and much more easily and probably successfully. It must be emphasized that this approach presumes the availability of reliable dimensions, and furthermore one has little insight into the inherent errors of such published dimensions when they are available.

The far more common case is to encounter debris as the secondary object, and such objects almost never have known dimensions. The only available signature information about the object is the average radar cross-section (RCS) value presented in the CDM. Were a full distribution of radar hit-level RCS values available for the object, it perhaps would be possible to pursue an object-specific size estimation activity; as it is, one is forced to make assumptions and rely on canonical distributions, with the averaged RCS value the only object-specific datum.

A study some years ago examined full hit-level RCS histories for all catalogued objects to determine which canonical distributions best describe the actual RCS PDFs for different object types.¹³ Examined were the set of Swerling RCS distributions, originally developed for aircraft, that are applied to space objects in the planning software of most space surveillance radars, as well as other distributions suggested in the literature, such as the lognormal distribution.¹⁴ It was found that the Swerling I and Swerling III models (Gamma distributions with shape parameters 1 and 2, respectively), along with a “middle” Swerling model consisting of a Gamma distribution with shape parameter 1.5, could reasonably represent the PDFs of about 62% of the catalogued debris objects. Of these three models, the Swerling III is the one that includes the largest objects and therefore of the three is the most conservative choice; and given how little is actually known about sizes of catalogued debris, a conservative selection would seem appropriate. The scale parameter for a two-parameter Gamma distribution is estimated by Maximum Likelihood methods as the sample mean divided by the shape parameter. The RCS value provided on the CDM can arise from a number of different calculation approaches, depending on the reporting sensor; a conservative interpretation is to presume that it represents a median value and therefore must be converted to a mean value for a Swerling III distribution; this is done by multiplying the median value by a factor of 1.19 (a mean-to-median ratio that can be obtained by a straightforward Monte Carlo generation of the distribution). With the assumptions given above, one can use the RCS value provided in the CDM to construct a properly-scaled PDF of the expected RCS distribution for the debris object.

To be useful to the present enterprise, of course, this RCS distribution must be converted to a size distribution. Fortunately, an experiment was conducted by the NASA Orbital Debris Program Office that allowed the development of an algorithm to do this.¹⁵ A test satellite was exploded in a vacuum chamber, the pieces preserved, and thirty-nine representative pieces observed (in all orientations) by a small radar in an anechoic chamber over a large range of radar frequencies. From these data a functional relationship was defined between the resultant RCS and the object “size” or, more correctly, characteristic dimension (defined as the average of three orthogonal vectors inscribed in the object in a certain way). This relationship has certain restrictions on its use: it is intended only for debris objects and even then those smaller than 20 cm, and it is legitimately used only to convert a distribution of RCS values to a distribution of expected sizes (not to convert a single average RCS value to a single resultant average size). While the relation-

ship is best applied to measurements taken on a large group of objects, one does not stray far from the original intent in applying it to a constructed distribution of RCS values for a small debris object.

The procedure and its results can be illustrated through the graphs in Figure 5. The average RCS value for the object is 0.02 m^2 , and using this value to calculate the needed scale parameter produces the Swerling III RCS distribution in the left pane. This distribution is then converted by the use of the above-described size estimation model to a characteristic size distribution, which is shown in the right pane. As can be seen from the actual values produced, if the primary object is of any reasonable size, the contribution of a small debris secondary object to the overall HBR will be very small.

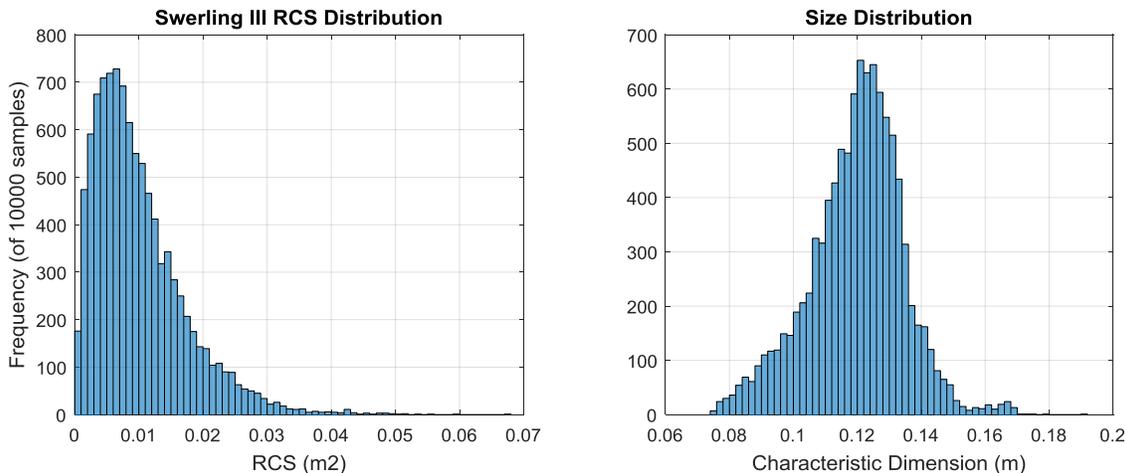


Figure 5. RCS Swerling III RCS and Size Distributions for Satellite with Average RCS of 0.02 m^2 .

PC CALCULATION NATURAL VARIABILITY

As mentioned previously, the two-dimensional P_c calculation yields a single point estimate (essentially an averaged value) and, as presently constructed, does not give any confidence interval information or sense of the underlying distribution of which that point estimate is a representation. This fact has not escaped the notice of conjunction assessment practitioners, and one attempt to compensate for this is to calculate and examine the distribution of miss distance values. It is straightforward and not particularly computationally burdensome to generate this information with a simple Monte Carlo approach: one can perform a random draw on the covariances propagated to the time of closest approach (TCA) for the nominal solution, add the drawn perturbations to the nominal solutions, and, using rectilinear motion (this is a reasonable approximation for small perturbations within the conjunction region), determine the new TCA for this pair of positions and velocities and the miss distance that at that new TCA. Executing a large number of repetitions of this procedure will produce the desired distribution, which one can then examine as either a probability density function or a cumulative distribution function. In fact, one need not even resort to Monte Carlo, as a direct analytical solution for the miss distance distribution was derived by Chan¹⁶ and demonstrated in practical use by Coppola *et al.*¹⁷ While these data can be produced easily enough, the difficulty arises in how precisely to interpret them for operational purposes. One can calculate the fraction of the miss distances that fall within the hard-body radius and thus determine a collision probability, but this is just the computation of the P_c through a different means and thus a return to a single parameter. One could examine the distribution of the

miss distances in the neighborhood of the HBR and thus understand how the Pc would vary if the HBR were altered slightly, but this is more a sensitivity investigation of the HBR and not of the miss vector distribution proper. Finally, one could try to interpret the full presentation of the miss vector distribution as a PDF or cumulative distribution function (CDF), but it is difficult to draw operationally-useful conclusions from this since it is not framed as a parameter that relates directly to risk.

An approach that has presented itself for experimentation is to choose a random position draw from the combined covariance and, using the miss distance that this produces as the nominal miss distance, to proceed to calculate the Pc with this new nominal miss and the unaltered combined covariance. A large number of such repetitions would produce a collection of Pc values, whose distribution could then be characterized. A resampling technique of this type for the present situation is not an unreasonable construct, but the approach as formulated here is not tenable because the combined covariance is correlated to the nominal miss vector, in terms of both size and orientation. One cannot generate a new nominal miss vector from sampling from the combined covariance and then use, unmodified, the combined covariance as a reasonable representation of the position uncertainty about that new nominal miss vector.

A clever modification by Frisbee¹⁸ of the resampling technique addresses this problem while at the same time bestowing additional benefits. Beginning with the nominal miss vector and combined covariance, m componentized samples are taken from the combined covariance. The mean of these samples is calculated and is combined with the “old” nominal miss to define a new nominal miss for this sample set. A sample covariance of the m samples is then calculated and serves as the combined covariance for this sample set. From this new nominal miss and covariance, a Pc can then be calculated with one of the many 2-D calculation techniques. Repeating this procedure n times will produce a distribution of Pc values that, it is proposed, can serve as a statement of the expected variability in the calculation.

This form of the resampling approach addresses the problem of correlation between nominal miss distance and covariance by ensuring that the covariance, which is a sample covariance calculated from the set of position draws from the original combined covariance, derives organically from the new nominal miss distance, which is the mean of these position samples. It does, however, introduce a new difficulty, namely determining the number of samples m that should be taken in each of the n repetitions of the resampling; and to answer this question one reflects on the orbit determination (OD) process that produced the position estimates for the primary and secondary states. OD is an estimate of a mean state using a certain amount of sampling data from the actual orbit (*i.e.*, observations), and the number of degrees of freedom (DoF) of the estimate is the difference between the total number of observables used in constructing the OD and the number of parameters estimated. If a resampling procedure is used to determine confidence intervals for calculations that derive from a particular OD, the resampling should be configured so as to employ the same DoF as the OD to which it is related. This would align with standard bootstrapping rule that, to determine confidence intervals about a parameter estimated from n samples, the additional datasets to be drawn (with replacement) for the resampling should also all be of size n .

The concept becomes somewhat muddled when applied to relative miss distances and combined covariances, as they are the mix of two different OD solutions (one for the primary and one for the secondary), each of which may have very different levels of tracking. Cleanliness, however, can be returned to the situation by applying the resampling approach independently to each object: m position samples are drawn from the primary object’s covariance and a mean miss distance and sample covariance generated; the same is done for the secondary object, although the number of samples drawn may be different from m (to be discussed shortly); the two miss dis-

tances are used to construct a new nominal miss, the two covariance matrices are combined into the single combined covariance, a 2-D Pc is calculated in the usual way, and the procedure is repeated n times to generate a Pc distribution. A diagram that illustrates this process is given in Figure 6 below:

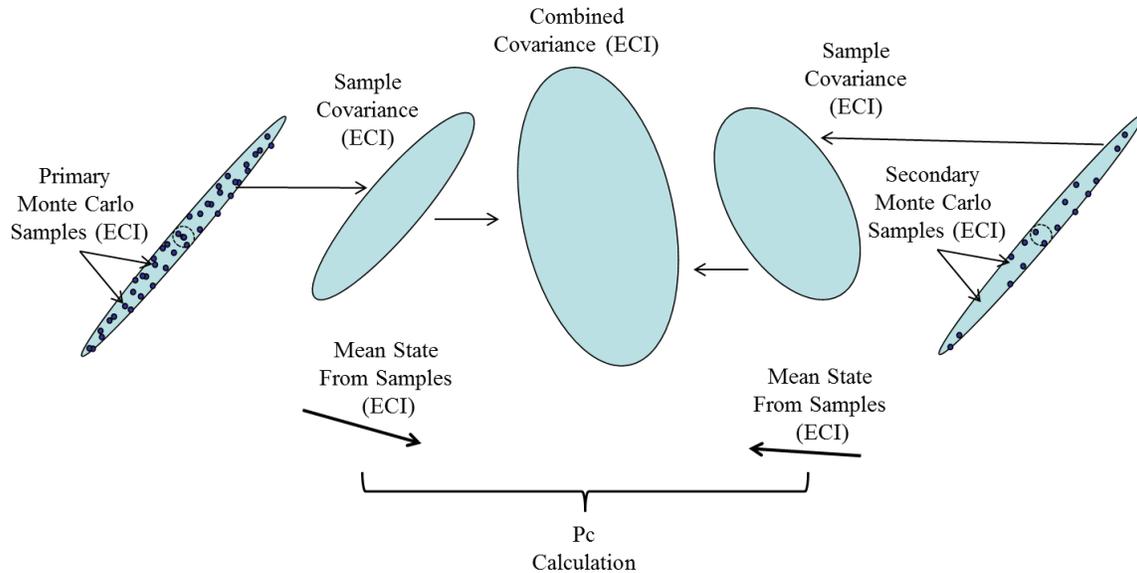


Figure 6. Pc Calculation Resampling Technique Schematic.

As discussed earlier, the number of samples taken from either the primary or secondary covariance to determine the mean state and sample covariance should be such that the DoF of those calculations should equal the DoF of the associated OD. But how should the DoF of a JSpOC OD be determined?

Space Surveillance Network radar observations, which constitute the great majority of the tracking data received and nearly all of the tracking data on near-earth objects (those with periods less than 225 minutes), typically each contain three observables: a range measurement and two angle measurements. These observations are usually furnished in “tracks” of observations, meaning sets of observations arising from the same satellite-sensor observing session; three to ten observations collected and submitted this way would be a common situation. One could conclude, therefore, that a track of ten observations would furnish 30 independent and useful observables; but the correlation among observations in the track (which are taken in close temporal proximity) is so large that these individual observations are hardly independent contributors to the OD. It is thus not clear how many independent elements truly exist in an observation track, especially because the individual observations themselves are the product of a tracking filter. To add yet more complication, the JSpOC OD process also presents some ambiguities in determining the number of parameters estimated in an OD. Because along-arc solutions through fit-span segmentation are pursued for the non-conservative forces (atmospheric drag and solar radiation pressure), which result in multiple ballistic coefficient and solar radiation pressure coefficient solutions, it is not obvious precisely how many independent parameters are being estimated; and in any event there is no record available to the user of a CDM that indicates the number of segmentations pursued, so it is not possible to account for this in the risk assessment process anyway.

Given these limitations, the following DoF estimation procedure is proposed. Rather than trying to identify the individual items of information content in a track of observations and similarly the individual parameters that are part of the state estimate in the OD process, one can manage the accounting by allowing the “state estimate” to be the unit of measure: the generated state estimate itself is one degree of freedom (one “state”), and the observational data will be quantized by the amount of such data that would allow a single state to be estimated. Given some of the limitations of and problems with Space Surveillance Network tracking information, how many such observations should be considered the equivalent data unit of the estimation of a single state?

If the track of observations contains several metric observations and each of them further contains three observables, then one can suppose that the track could contain about enough information to perform a state estimate update. To be sure, the update propriety for the non-conservative force values (*i.e.*, ballistic coefficient and solar radiation pressure coefficient) is likely to be unsatisfactory for a batch update process, but the group of observations should contain about the amount of information necessary to execute a reasonable OD. Similarly, rather than try to identify the total number of parameters estimated in an OD update in which along-arc methods are used, it seems acceptable simply to consider the OD output as the estimate of a single state—not multiple parameters (even though of course it is composed of that) but a single estimated state. Using this approach, the number of degrees of freedom in a particular OD update is calculated as the number of tracks of reasonable length minus one (the one “state” that is estimated). This process requires further work for situations in which angles-only tracks are common, as it is probably not the case that that a track of angles-only observational data really is sufficient for a state estimate; but it does seem a reasonable initial procedure for near-Earth orbits.

Further work is also needed to determine the number of repetitions of the procedure— n from the above discussion—needed to reach a stability condition. Preliminary efforts have indicated that in most cases stability is reached with 5000 trials; but adequate performance may frequently be possible with less, and such a result would bring a welcome reduction in the execution time.

OMNIBUS EXECUTION AND RESULTS

All three perturbation types described above (covariance realism, HBR, and Pc calculation natural variation) are folded into a nested Monte Carlo framework for omnibus execution. For the present, an equiprobable gridding approach is taken for the covariance realism values in order to improve execution time; thirty equiprobable bins are defined for the scale factor probability density for the primary and thirty for the secondary, the median value from within each bin is used as the representative value for that bin, and all combinations of these gridded values ($n \times n$) are applied to the primary and secondary covariances as part of the outer Monte Carlo loop. The resampling approach is then run for each pair of scaled covariances, presently with 5000 samples per invocation. The resampling results are then reduplicated to allow the assignment of a binned set of HBR values according to the joint probability density for the primary and secondary object sizes.

Each of the perturbations (covariance realism, HBR, and resampling) is also run by itself, with the other two inputs either not executed or set to the nominal value, to see what probability density is produced by the variation of that item alone; this reveals the sensitivity of the particular conjunction to the particular item. Reasonably well-tracked cases that nonetheless have schizophrenic covariance realism performance will, for example, show a large probability density spread for covariance realism perturbations but a relatively modest one for natural variation.

Figure 7 provides a proposed display of results. The left panel gives the four probability densities, shown here as cumulative distribution functions (CDFs) to allow the full range of probability points to be displayed in a comparative context; the nominal P_c value is represented by the magenta dotted line. The right panel redisplay this same information, this time giving 5th- to 95th-percentile error bars, with a large dot at the median value. While this panel gives less complete insight into the full probability density of each item, it does give a more immediate visual take-away of how each of the perturbation sets and the combined output compare to the nominal value. In the example given in the figure, natural variation of the calculation is more or less evenly balanced about the nominal value; but covariance realism considerations substantially push the risk almost entirely below the nominal value, and adding a more realistic HBR decreases the risk alone by about half an order of magnitude. The Total (blue) line, which collects all of the perturbations together, falls mostly below the nominal value. If a user's remediation threshold were a P_c value of 1E-04, one might hesitate somewhat at pursuing active remediation for this conjunction, as 70% of the total perturbation P_c values fall below the 1E-04 threshold.

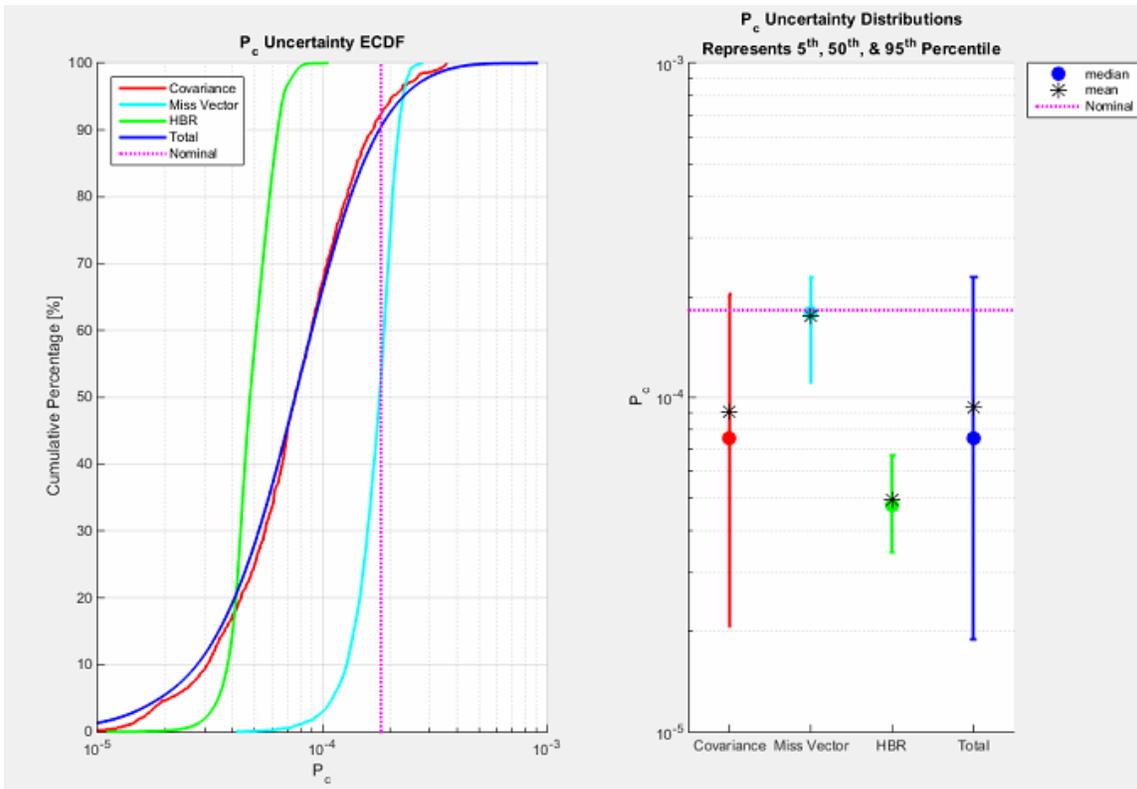


Figure 7. JAC CoPoC Sensitivity Display.

CONCLUSION AND FUTURE WORK

A construct has been defined by which the characterized uncertainties of the primary objects' covariances and HBR values, as well as the natural variation in the P_c calculation, can be used to generate a probability density of P_c outputs; this probabilistic output can be compared directly to P_c decision thresholds, allowing more informed decisions regarding the urgency of remediation activity.

The construct is embryonic and is just beginning to be explored with larger datasets and off-line operational exercise, and there are a number of areas to which additional analytical work will be applied. Most important is a more thorough examination and benchmarking of the overall Monte Carlo technique, which as presently formulated includes certain simplifications in order to improve overall performance; comparison and stability studies need to be conducted in order to determine how much error the simplifications introduce, how many samples and bin divisions are required in order to meet certain accuracy criteria, *etc.* More study will also be directed at the degree-of-freedom accounting approach used in the resampling technique, especially to extend the construct to angles-only observation tracks. Finally, while the construct is not, as previously explained, intended to serve as a bounded predictor of future Pc updates, it is of interest how well it will actually serve this role; so an examination of this kind of performance against historical data will be quite interesting and valuable.

ACKNOWLEDGMENTS

The authors wish to thank Mr. Richard Grist for the initial work on Pc uncertainty that has culminated in the present construct, Mr. Stephen Casali for a number of helpful suggestions regarding covariance realism parameterization and calculations, and Mr. Joseph Frisbee for multiple useful discussions about Pc calculation resampling techniques.

REFERENCES

- ¹ Foster, J.L. and Estes, H.S. "A Parametric Analysis of Orbital Debris Collision Probability and Maneuver Rate for Space Vehicles." NASA/JSC-25898, August 1992.
- ² A useful survey of many of these methods is provided in Alfano, S.A. "Review of Conjunction Probability Methods for Short-Term Encounters." AAS/AIAA Paper No. 07-148, AAS/AIAA Space Flight Mechanics Meeting, Sedona AZ, February 2007.
- ³ This phenomenon was first described in detail by Alfano, S.A. "Relating Position Uncertainty to Maximum Conjunction Probability." AAS/AIAA Paper No. 03-548, AAS/AIAA Astrodynamics Specialists Conference, Big Sky MT, August 2003.
- ⁴ Laporte, F. "JAC Software, Solving Conjunction Assessment Issues." 2014 AMOS Technical Conference, Maui HI, September 2014.
- ⁵ Vallado, D.A. and Seago, J.H. "Covariance Realism." AAS/AIAA Paper No. 09-304, AAS/AIAA Astrodynamics Specialist Conference, Pittsburg PA, August 2009.
- ⁶ Cerven, W.T. "Improved Empirical Covariance Estimation." AAS/AIAA Paper No. 13-762, AAS/AIAA Astrodynamics Specialists Conference, Hilton Head SC, August 2013.
- ⁷ Hornwood, J.T., Aristoff, J.M., Singh, N., Poore, A.B., and Hejduk, M.D. "Beyond Covariance Realism: a new Metric for Uncertainty Realism." Proc. SPIE 9092, Signal and Data Processing of Small Targets, 2014 90920F (June 13, 2014).
- ⁸ Hejduk, M.D., Ericson, N.L., and Casali, S.J. "Beyond Covariance: A New Accuracy Assessment Approach for the ISPCS Precision Satellite Catalogue." 2006 MIT / Lincoln Laboratory Space Control Conference, Bedford, MA. May 2006.
- ⁹ For a conceptual development of this topic, see Hejduk, M.D., Plakalovic, D., Newman, L.K., Ollivierre, J.C., Hametz, M.E., Beaver, B.A., and Thompson, R.C. "Trajectory Error and Covariance Realism for Launch COLA Operations." 2013 AAS/AIAA Space Flight Mechanics Meeting, Kauai, HI. February 2013.
- ¹⁰ Chan, F.K. *Spacecraft Collision Probability*. El Segundo, CA: The Aerospace Press, 2008. The relationship follows directly from Equation 5.28.

¹¹ A number of examples of different sphere-circumscription approaches chosen by satellite owner/operators are given in Vincent, M.A. “Bridging the Gap between Academia and Operations for Orbital Debris Risk Mitigation.” 2015 AMOS Technical Conference, Maui HI, August 2015.

¹² Matney, M. “How to Calculate the Average Cross-Sectional Area.” *Orbital Debris Quarterly Newsletter*, Vol. 8 No. 2 (April 2004).

¹³ Hejduk, M.D. and DePalma, D. “Comprehensive Radar Cross-Section ‘Target Typing’ Investigation for Spacecraft.” 2010 AAS/AIAA Astrodynamics Specialists Conference, Toronto CA, August 2010.

¹⁴ Seybold, J.S. and Weeks, K.L. “Arithmetic versus Geometric Mean of Target Radar Cross-Section.” *Microwave and Optical Technology Letters*, Vol. 11, No. 5 (April 1996), pp. 265-270.

¹⁵ The basic construction and use of the developed model is discussed in Barton, D.K. *et al.* “Final Report of the Haystack Orbital Debris Data Review Panel.” NASA Technical Memorandum 4809, February 1998. Additional information about its use is provided in Rajan, N. *et al.* “Orbital Debris Size Estimation from Radar Cross Section Measurements.” *Proceedings of the Third European Conference on Space Debris*. Darmstadt Germany, March 2001.

¹⁶ Chan, F.K. “Miss Distance – Generalized Variance Non-Central Chi Distribution” AAS/AIAA Paper No. 11-175, AAS/AIAA Astrodynamics Specialist Conference, Girdwood AK, August 2011.

¹⁷ Coppola, V.T., Woodburn, J., and Hujsak, R. “Effects of Cross-Correlated Covariance on Spacecraft Collision Probability.” AAS/AIAA Paper No. 04-181, AAS/AIAA Space Flight Mechanics Conference, Maui HI, February 2004.

¹⁸ Frisbee, J.H. “International Space Station Collision Probability Analysis.” JSC Technical Memorandum OFD-03-48300-010, 2003.