SATELLITE CONJUNCTION "PROBABILITY," "POSSIBILITY," AND "PLAUSIBILITY": A CATEGORIZATION OF COMPETING CONJUNCTION ASSESSMENT RISK ANALYSIS PARADIGMS

M.D. Hejduk[*] and D.E. Snow[†]

A number of different conjunction assessment (CA) risk analysis methods and metrics have been proposed in the critical literature, and they vary widely in purport and form. However, they tend to be proposed individually and episodically, so that it is difficult for a CA practitioner to take stock of the possibilities, understand their fundamental differences, and make informed choices for their particular CA risk assessment enterprise. The present study seeks to collect the major proposals for risk assessment methods and parameters and organize them categorically, under the proposed divisions of "probability," "plausibility," and "possibility," as well as formulate what appears for each to be its fundamental question and, where applicable, null hypothesis. This activity can, through a bottom-up approach, provide some of the building blocks for an overarching CA philosophy, as well as establish concepts and terminology potentially useful to the broader discussion of these topics.

INTRODUCTION

A number of different conjunction assessment (CA) risk analysis methodologies have been proposed in the critical literature. Some of these are straightforward and have been used operationally for some time, some are new and novel formulations, and some are repairs and expansions that have been offered in response to perceived deficiencies in existing approaches. Regardless of content and origin, however, such methods and metrics tend to be proposed episodically and in a stand-alone manner; and the lack of a common vocabulary to describe them, or a common framework within which to place them comparatively, has made it difficult for CA operators to make informed choices among them when setting up their CA enterprise or working to evolve it towards greater sophistication. To provide the conceptual background for informed risk analysis metric choices, a full philosophy of CA needs to be developed, with all of the different motivations and considerations recognized and placed within a categorical system. Some work has begun on this score, although it has become a moving target due to all of the recent and rapidly-unfolding changes in the population and use of earth orbits. However, despite the present lack of such a top-down treatment of CA, one can nonetheless provide useful information and assistance through a bottoms-up approach of considering the major risk analysis metrics that have been proposed and offering a categorical scheme that will help to make evident their similarities

---

[*] Chief Engineer, Astrorum Consulting LLC, 10006 Willow Bend Drive, Woodway TX 76712
[†] Senior Scientist, Omitron Corporation, 555 East Pikes Peak Avenue (#205), Colorado Springs CO 80903

and differences. Such a system of categories is not at all a substitute for a supervening philosophy of CA, but establishing a taxonomy of the major CA risk analysis metrics that have been proposed to date can certainly illuminate aspects of the risk analysis enterprise that can help guide the assembly of this overall philosophy. Establishing a common vocabulary with which to discuss different risk analysis approaches, both those proposed to date and, potentially, future proposals, could also be very useful.

This study will begin by proposing a categorical grouping that is believed can distinguish and arrange the different risk methods and parameters most usefully: the degree to which the approach appeals to probability versus possibility of collision. After treating this principal distinguishing feature, additional general attributes of these risk analysis methods will be outlined; while they do not always serve as robust categorization principles, they do usefully serve to tease out important attributes of the approaches and provide additional insight into their differences. Following this outlining, a number of proposed risk analysis parameters from the literature will be individually summarized, discussed, and categorized. The purpose of this latter activity, which actually is the focus of this study, is not to evaluate the propriety of any particular risk analysis parameter or method; rather, it is to show conceptually where the parameter falls in a categorization scheme and therefore to provide the initial ingredients for determining the situations for which the parameter is best suited. It is hoped that the proposal of a categorical scheme and the activity of and discussion regarding the placement of each technique into this scheme may be helpful to CA practitioners wishing to evaluate which proposed techniques may fit the particular philosophy and objectives of their enterprise.

LINGUISTIC MOOD AND SATELLITE CONJUNCTION RISK ASSESSMENT

Of the different linguistic moods that exist in Indo-European languages, there are three that address the degree of certainty of the described event or circumstance: the indicative mood, which is an expression modality of certainty and used to assert that a particular circumstance will occur; the subjunctive mood, which is a modality of less certainty and expresses that a particular circumstance may occur; and the optative mood, which is a modality of potentiality and expresses what could occur. It is not an accident that three distinct linguistic forms developed to discuss different levels of certainty; and in surveying the different CA risk analysis parameters and methods, one begins to see an alignment of each approach with one of these three linguistic moods in that some attempt a precise calculation of collision likelihood (indicative mood), some attempt to consider additional errors in the process and therefore speak essentially in terms of likelihood of likelihoods (subjunctive mood), and some focus simply on establishing or negating the possibility of a collision event (optative mood). Because these linguistic concepts have essentially become terms of historical grammar, given that most modern languages lack an optative mood and grammatical modernists are doing their best to eliminate the English subjunctive mood, it may be better to try to establish a more direct vocabulary. To this end, the terms probability, possibility, and plausibility are proposed.

These terms are introduced with some reluctance because, while not as remote as Indo-European grammatical concepts, they may introduce confusion and misunderstanding more than they promote clarity, for a couple of reasons. First, risk assessment parameters do not neatly fit into three bins as defined by these three concepts; rather, there is more of a spectrum that runs from an (ostensible) precise statement of collision risk to an assessment of mere collision possibility. Second, the difference in meaning between possibility and plausibility is not clearly defined and takes on a somewhat idiosyncratic form for this analysis, as well as potentially creating dissonance with both more precise meanings of the terms within statistical science and previous uses with regard to CA concepts. Nonetheless, these particular terms do seem to be functional discriminators for the current categorization task, so they will be proposed here despite certain reservations with the following definitions.

Probability

The risk assessment parameter is a determination, or at least an attempt at a determination, of an actual probability that a serious close approach event will take place. There are disputes in the literature regarding whether some of these parameters that purport to be probabilities truly are[1,2]; but since it is the case that many practitioners believe these parameters to represent probabilities and treat them as such, it makes sense that they be categorized this way. These parameters can take the form of a single point-estimate of probability for a given collection of conjunction data or a broader statistical structure that is still probabilistic in form.

Possibility

The risk assessment construct seeks to determine whether a collision is simply possible, to a certain level of confidence, with the idea that one would act to improve the safety posture based merely on the possibility of a collision. There are certainly statistical concepts involved in this determination, as some sort of threshold must be established to determine whether there is less than a negligible opportunity for collision and therefore no appreciable possibility; but there is no outright calculation of collision likelihood.

Plausibility

"Plausibility" is the term suggested here to stand between probability and possibility and thus describe approaches that are not simply an attempted calculation of the likelihood of collision yet do not focus on the mere possibility of collision. Such approaches often take the form of a probabilistic calculation but attempt to consider additional uncertainties regarding the input data and/or the overall problem. This can be done in a largely to entirely quantified manner, in which the approach could properly be considered an expansion of the "probability" category; or it can be more speculative and consider intuitive ranges or potential maxima. Overall, the attempt is to determine, even in the wake of a direct probability calculation that would generally not be considered worrisome, whether a serious situation is still plausible (but not merely "possible") and therefore whether a mitigation action should be considered.

THE RISK ANALYSIS FUNDAMENTAL QUESTION

All decision processes are, or at least should be, guided by an answerable and focused question that the decision process is attempting to answer. One could say that the question that illuminates the key CA issue is "Will the two objects collide?" While perhaps correct in being the foundational question of CA risk assessment, this question requires appropriating the "God's eye" truth of the situation, which is not actually possible. In the absence of such access, one has no choice but to proceed statistically and thus needs a more focused form of the question, namely a statistical one. In the above statement, the adjective "focused" is important—the question must be neither too broad nor too specific or else it will be ineffective at guiding the decision process. For example, promoting the safety of an individual satellite mission or preserving a particular orbital corridor from excessive debris pollution are both important goals of the CA process, but these concepts are too abstract to be illuminating to a particular risk analysis decision for a particular satellite conjunction. As such, a question of the form "Would a satellite maneuver away from a collision to promote the safety of its occupied orbit regime?" does not really assist the decision nexus of a potentially worrisome conjunction because any improvement in safety, no matter how minor or expensive, would yield an affirmative response. Similarly, a question that is extremely specific and technical might be an actual question that is asked as part of the decision process and contribute helpfully to the risk analysis conclusion, but it is not a fundamental question because it does not provide basic guidance to the decision process.

These initial comments seem almost to be suggesting use of the *via negativa*, by which one establishes what a thing is by stating that which it is not. There are, however, some attributes of the fundamental question that can be directly specified. First, while the risk analysis process does attempt to elucidate a number of aspects of the conjunction, the overriding consideration is usually whether or not a mitigation action to lessen the collision risk should be pursued; this is what propels risk analysis beyond mere situational awareness to planning and ultimately ruling on the desirability of possible mitigation activities. The fundamental question should address this most meaningful aspect of risk analysis—should the conjunction's risk be mitigated?

Second, the fundamental question should also establish a default position with respect to mitigation actions when the risk assessment data do not allow a fully certain or convincing conclusion. In decision theory this issue is called *presumption* because it indicates the "presumed" action (or lack of action) that will be taken unless the evidence is convincing in favor of a change in course; it thus indicates the side on which the burden of proof falls and how one will act should this burden of proof not be met. Certain risk analysis parameters encounter situations in which they are not fully representative of the risk attribute they are attempting to represent, and the fundamental question should make clear how one is to behave in such circumstances. In fact, perhaps the most contentious and important issue in conjunction risk analysis presently is whether CA operators should favor or refrain from mitigation actions in the presence of ambiguous data.

THE NULL HYPOTHESIS

The way CA risk analysis is typically conducted, regardless of whether the practitioners themselves consciously conceive of it in these terms, bears a certain outward similarity to statistical hypothesis testing. The risk analysis process employs a conjunction likelihood or similarly-functioning test statistic; the probability of collision is the test statistic most frequently used, but a number of alternative statistics have been proposed. The test statistic is then compared to a threshold, which will define a critical region in the response space; if the test statistic is in the critical region, a particular action is nominally warranted. The further the test statistic penetrates into this critical region, the greater the force of the conclusion that this particular action should be performed. Confidence intervals can be defined for the test statistic, which will allow the potential penetration into the critical region to be stated with further precision. All of these aspects of the decision process manifest a strong similarity to classical hypothesis testing, in which a null hypothesis is stated, a test statistic is established, a critical region for the test statistic is also defined, and a particular conclusion is embraced if the test statistic is located in the critical region. The full degree of similarity of the CA risk analysis process to a hypothesis testing situation is an issue that is currently under debate in the literature[1,2], just like the entire concept of hypothesis testing is itself under debate in the statistics community[3]; and a more definitive resolution of the precise statistical posture of this process is an important issue to be resolved for the CA community. For the purposes of the present analysis, the goal of which is to attempt a taxonomy of the major proposed risk analysis parameters, it is sufficient to maintain that the risk analysis operates in a similar way to hypothesis testing, namely that there is a null hypothesis of sorts defined, a test statistic established, a critical region for the test statistic defined, and a rejection of the null hypothesis should the test statistic be placed in the critical region. As previously implied, it is not at all clear that CA practitioners commonly conceive of their risk analysis activities in these terms; but in the authors' experience it is a robust model for describing how most CA risk analysis is actually performed. It therefore seems reasonable to employ the concept of the null hypothesis as a potential taxonomic principle for categorizing risk assessment parameters.

EVENT DATA

An additional method to distinguish among different risk analysis test statistics, although it is more descriptive than foundational, is the conjunction event data that they require for their computation. The Conjunction Data Message (CDM), which is a Consultative Committee for Space Data Standards (CCSDS)-defined message for reporting relevant data related to conjunction assessment, contains no small amount of useful information; principal among its contents are the states and covariances of both of the objects at the time of closest approach (TCA), radar cross-section information, tracking information, and orbit determination (OD) force model information. Many risk analysis parameters can be calculated directly from the information given in the CDM. Some techniques do not strictly require additional data beyond the CDM's contents but do make certain assumptions about the character and behaviors of some of these data. Other techniques require *a priori* parameters and statistics that are the product of external analysis. As stated above, the consideration of type of input data, and the "distance" the input data are from what is directly available from the CDM, do not constitute principles of categorization *per se*; but they are both quite useful as descriptors and do play a small philosophical role in that they help to define whether the risk assessment methodology itself treats the conjunction as truly a stand-alone event or tries to inform the risk analysis situation with information from past data histories.

DESCRIPTION OF DIFFERENT RISK ANALYSIS PARAMETERS/TECHNIQUES TO ESTABLISH CATEGORY

Probability of Collision Calculation

The probability of collision (Pc) calculation, whether based on a 2-D analytic calculation,[4,5,6,7] a 3-D analytic calculation,[8] or Monte Carlo techniques,[9,10] is the most widely used CA risk analysis parameter. In considering both the state estimates and uncertainties of the primary and secondary objects at TCA, this parameter represents the likelihood that the relative miss vector will be smaller than a user-defined hard-body radius (HBR), or range of hard-body radii, taken to represent the satellites' combined sizes. Usually the test statistic threshold that will here define the critical region is relatively small; most CA practitioners who employ this technique (which is at present the great majority) choose a threshold in the neighborhood of 1E-04, meaning that they will typically counsel a mitigation action if the probability of a collision exceeds one in ten thousand. All of the Pc calculation techniques listed above can be executed with data provided on the CDM, and in fact all that is truly needed are the two objects' states and covariances at TCA, as well as some means by which to determine a value for the HBR.

The Pc calculation exhibits a peculiar phenomenology called the "dilution region,[11]" namely that a low-Pc value can arise from either a very precise knowledge of the two objects' positions (thus concluding definitively that they are unlikely to collide) or an extremely coarse knowledge of the positions (because the possibilities of future positions are so broad, one cannot conclude that a collision is likely). The practical effect of this phenomenon is that, while high-Pc values can be seen as evidence of a serious conjunction situation, low-Pc values do not allow a conclusion that the situation is safe.

If the Pc is the only risk analysis parameter employed, the dilution region phenomenon more or less of necessity is handled operationally in the following way: to counsel remediation when the chosen Pc threshold is exceeded and to refrain from a remediation action as a default condition when this threshold is not reached. A low Pc means that one cannot conclude that a worrisome CA situation is present, so no action is taken (the default condition); a high Pc means that there is good evidence of the presence of a worrisome condition, so one can justify a mitigation action. This is not an inherently unsafe posture because it is aligning the events with data insufficiently precise to allow a conclusion of high risk (that is, dilution region events) with conjunction events with objects too small to be tracked and thus not in the satellite catalogue at all (a limiting case of data imprecision—uncatalogued objects are essentially the same as objects with a random state estimate and infinite covariance). Dilution-region events are thus seen as part of the overall background risk of collision with objects about which too little is known to mitigate definitively. This position is developed in more detail in Hejduk *et al.*[12].

Embracing this position, which requires a default posture of non-mitigation, suggests the following fundamental question for the CA decision: "Do the presented data justify a decision to mitigate the conjunction?" The question in this form both focuses the decision on the potential mitigation action and properly identifies the default situation. In seeing the decision nexus as analogous to hypothesis testing, the null hypothesis would thus take the following form: "The actual miss distance is greater than the hard-body radius." This is the default position that would be rejected should the Pc value exceed the defined threshold.

It should also be stated that, if the use of the Pc were to be combined with other techniques, then the fundamental question and null hypothesis might be expanded or become multi-part. For example, the Pc might be used as the deciding parameter for non-dilution-region situations, and dilution-region events might be allocated to a different risk analysis technique. But for a construct that uses the Pc alone as the risk analysis metric, the fundamental question and null hypothesis given above seem the most compelling.

Wald Sequential Probability Ratio Test

Developed by Carpenter and Markley,[13] the Wald sequential probability ratio test (WSPRT) is an interesting and impressive proposal that accounts for all aspects of the CA decision environment: it provides a test statistic (the Wald sequential probability ratio), a construct that allows sequentially-generated conjunction data to be processed and accommodated, a transformation paradigm to generate alarm (should maneuver/mitigate) and dismissal (no action required) thresholds for the test statistic based on user-specified false alarm and missed detection levels, and the familiarity of working within a Pc-based framework. Its framing principle is to determine how much worse the current conjunction is from what is presented by the usual interaction of the two objects over time, and this concept is what gives form to the test statistic itself, which is the ratio of the collision likelihood of the current event to the collision likelihood between the two object in general. This ratio is calculated sequentially as CDM updates are received, and the time-series data can be displayed in the following way:

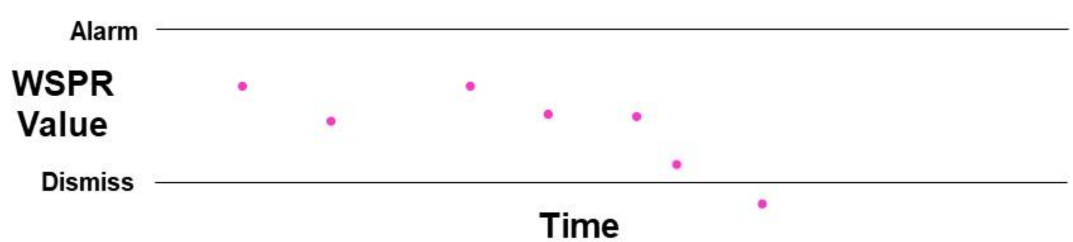

Figure 1: Time-series Wald Sequential Probability Ratio Test (WSPRT) data, compared to alarm and dismissal condition thresholds

If the alarm threshold (the top line in Figure 1) is violated, then the combination of false alarm / missed detection criteria indicate that a mitigation action should be pursued, and one can act on this immediately if desired. Similarly, if the dismissal threshold (lower line) is breached (last pink point in Figure 1), then the event can be dismissed at that point. If the value stands between the two thresholds, then the test is indeterminate; and one should wait for additional data until one of the two thresholds is crossed.

Employing the WSPRT is fairly straightforward because the *in situ* data that it requires are all derivable from the CDM. However, the ratio also requires a background risk value between the two objects in conjunction ($P_{c|o}$), which is the usual conjunction risk that they pose to each other apart from this particular conjunction. This is an external datum that must be generated through analysis, and its calculation is not always simple or singularly defined.

Because the WSPRT very nicely has an independent threshold for each of the two basic actions considered for risk analysis (*viz.*, to mitigate or to take no particular action), one could say that no framing fundamental question is needed here: the construct itself focuses the decision on the mitigation/dismissal operational imperative, and the three states of the test statistic do not require a presumption or default condition. Or such would be the case if one always had the luxury of being able to wait until the test resolved itself by violating one of the two boundaries; but CA practitioners must make a mitigation decision at the maneuver commitment point, which is the point at which a final maneuver decision must be made in order for it to be finalized, commanded, and executed before TCA. If the WSPRT has not violated either boundary at that point, then the test result is ambiguous; and a default condition must be articulated.

An entire paper has been written on operational experience with the WSPRT,[14] which gives some insight into actual operational use of the construct. The WSPRT was developed for the intra-constellation CA among the four satellites for the NASA Multiscale Magnetosphere Mission (MMS); and in this mode, the default condition is to pursue a maneuver if practicable. The construct was also used as an advisory calculation for conjunctions between any of these four satellites and other catalogued objects, and in this context the default condition followed that for the Pc calculation, namely to pursue a maneuver only if the alarm condition is reached (that is, to retain the current trajectory if between the two WSPRT boundary conditions).

The WSPRT can be placed in the probabilistic category, but it is clear both from theoretical considerations and actual operational example that the fundamental question and null hypothesis for this approach can be framed to support a default condition of either mitigation or refraining from action. When such a selection is made, however, the test sustains a *de facto* conversion from a two-boundary test to a single-boundary test because the middle section between the two official thresholds is now effectively assigned to one of the two conditions. For the present categorization purpose, the test will be assigned a default position of refraining from remediation, as this is the posture presently assumed for its use in CA against the entire satellite catalogue.

CARA Pc Uncertainty

It has been recognized for some time that 18 SPCS-produced covariances are not always fully and properly representative of the actual state errors. While there have not been any formal studies published on this (given the sensitive nature of such a subject), covariance consistency checks conducted during individual conjunction events reveal that covariances earlier in the event often do not statistically represent states that occur at later points. The representativeness of the 18 SPCS covariances can be studied formally using an operational product called SuperCODAC, which creates a reference orbit for every catalogued object and generates position residuals and propagated covariances for every 18 SPCS-produced vector at propagation points of interest.[15,16] Such a dataset can be used to create, for every object for key propagation points, PDFs of scale factors that were historically necessary to make that object's covariance achieve the realism condition of producing Mahalanobis distance products that conform to a three-degree-of-freedom chi-square distribution. These scale factors can be used in connection with (usually) the secondary object's covariance to produce a family of plausible covariances for this secondary object, and this set of covariances can be used with a Pc calculation utility to produce a PDF of Pc values for the conjunction.[17] Figure 2 shows potential output from a utility such as this; a CDF of the Pc values is presented, and a mitigation action is deemed appropriate if a certain portion of the Pc probability density falls to the right of a threshold value.

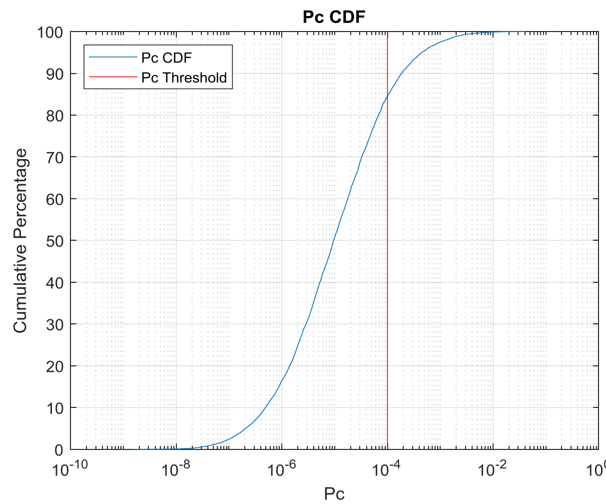

Figure 2: CDF of Pc Values from Pc Uncertainty utility

This method of presenting results is more complicated than that for a simple Pc calculation because one must specify both a Pc threshold (red line) and a second level (*e.g.*, 15%) of which the probability density exceeding the threshold must be greater in order to require a mitigation action. A somewhat more satisfying alternative is perhaps to use the inputs to each of the calculated Pc values in the above distribution to create a set of miss distance samples and then simply plot the CDF of these samples: by reading off the cumulative probability for the point at which the CDF curve crosses the hard-body radius miss distance level, one can infer the equivalent of a Pc value. A conceptual sample of what this display might look like is given in Figure 3, with the red line representing the HBR level of 20m:

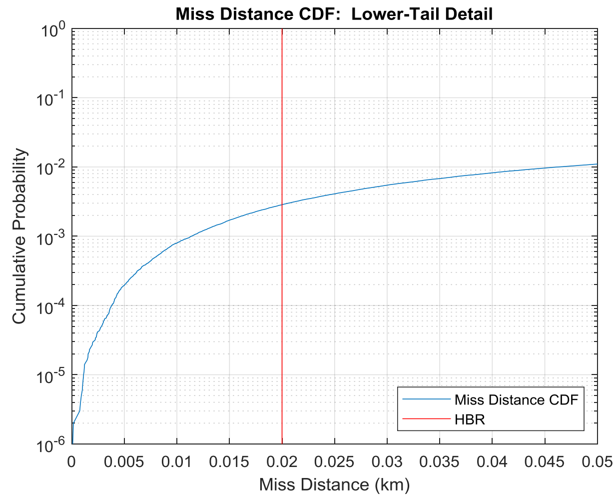

Figure 3: CDF of miss distances from Pc Uncertainty construct

The Pc Uncertainty approach takes some steps into a plausibilistic construct because it is asking what the plausible range of Pc (or miss distance) values might be if the current covariance were to suffer from the different levels of irrealism that have been observed in the past with the particular secondary object. However, it is noteworthy that the set of scale factors to apply to the secondary's covariance is not merely a range of imaginable values but an actual PDF, derived from a statistical covariance realism analysis of historical covariances for the object. Because a PDF of secondary covariances is used, a properly-weighted PDF of resultant Pcs is produced; and therefore the CDF's percentile points actually represent reasonable levels of expectation for the resultant Pc, under the assumption that the current covariance's history of irrealism is likely to be descriptive of the current covariance. This construct thus operates very much like the vanilla Pc calculation, except broadened to consider potential covariance errors.

The fundamental question recognizes this pedigree and is a slight broadening of that for the Pc taken alone: Given the current data and historical covariance realism information, does the Pc range of values justify a decision to mitigate? Since the data are to be used in a similar manner to that for the vanilla Pc, the null hypothesis will also favor maintaining the current trajectory: The actual miss distance is greater than the HBR.

Pc Sensitivity / Covariance Modulation

The Pc Uncertainty method described in the previous section had the luxury of a copious amount of historical covariance realism information computed against actual satellite reference orbits, so it was able to proceed in a statistically rigorous manner. Many CA practitioners do recognize that realism issues exist with 18 SPCS covariances and wish to compensate for them in some way, but they lack the historical realism information that would allow PDFs of scale factors or some similar artifice to be constructed. Some operators have thus pursued a boundary-value approach in which through experience they establish a minimum and maximum scale factor that they believe will encompass nearly all of the expected covariance variation and then assume a uniform distribution of values between these boundaries. If one proceeds this way for both the primary and secondary objects, it is possible to construct a grid of Pc values, each of which is a "plausible" Pc value given these assumptions. Figure 4 provides output from a version of this approach that Centre National d'Etudes Spatiales (CNES) has implemented.[18] The x-axis in the figure represents the scale factor applied to the primary's covariance matrix (range of values of

0.5 to 3), the y-axis the scale factor applied to the secondary's covariance matrix (again a range of values of 0.5 to 3), and the color the resultant Pc.

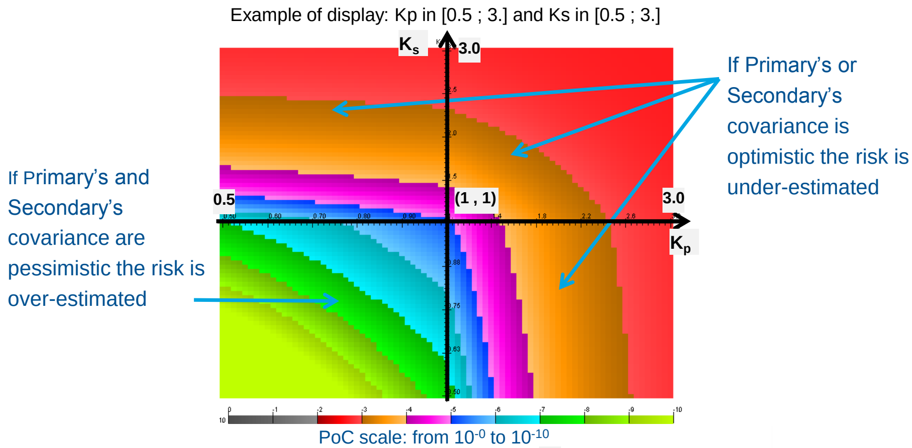

Figure 4: Intensity map of Pc as a function of primary and secondary covariance scale factor

In their operational implementation of this technique, CNES raised their Pc threshold from the value they had used with single-Pc calculations in order to make the threshold somewhat more lenient (*e.g.,* 1E-04 to 5E-04); but they now require a mitigation action if any of the calculated Pc values in the entire trade-space (*i.e.,* the entire colored grid shown in Figure 4) exceed this somewhat more lenient limit. This modification of operational procedure, in which CNES now recommends a mitigation action if any of these hypothesized Pc values exceeds their stated threshold, confirms that CNES has shifted their paradigm from a probabilistic to a more plausibilistic orientation.

This type of covariance modulation is more strongly plausibilistic than the Pc Uncertainty technique because the Pc values calculated for the trade-space are not actual probabilities; instead, they are calculations of what the Pc would be if the true covariances were actually mis-sized by the amount that each hypothesized ordered pair of scale factors indicates. To turn each of these into true probabilities, one would have to multiply each one by the likelihood that the two covariances would actually be mis-sized by the supposed amount. But they do represent plausible values of the Pc because they are derived from scale factors that fall between boundaries established from experience and limited experiment. The fundamental question is similar to that for Pc Uncertainty: Given the current data and covariance realism assumptions, does the Pc range of values justify a decision to mitigate? Regarding the null hypothesis, however, there does not seem to be a construct that is inherently favored by the method. If forced to select something, perhaps one would fall back on the fact that most events to not require a mitigation action, so the null hypothesis, as the "usual" situation, would be not to mitigate and would be rejected by even one above-threshold Pc in the overall trade-space.

Maximum Pc Construct

While there are a number of informal approaches to a maximum Pc construct, including small modifications to the usual Pc calculation, the best and most formal development of a comprehensive maximum Pc approach was offered by Alfano.[11] Developed in an era when most CA activities lacked covariance information, this technique addressed situations in which the basic "shape" of the covariance, namely the "aspect ratio" or ratios among the three covariance axes, can be established, as well as situations in which nothing at all is known about the covariance. The first of these two situations is the more interesting in that it applies directly to the dilution region phenomenon, so it will receive the more extended treatment here.

As stated previously, the dilution region is a situation in which the Pc value is depressed to a low value not because of high certainty but rather because of large ambiguity in the state estimates; this ambiguity manifests itself in larger covariances than would be observed if the state prediction errors were lower. Because the Pc will be low both if the joint covariance is very small or very large, it must reach a maximum somewhere between these two states. Therefore, if one begins with a very large covariance, successively shrinks it, and calculates the Pc at a large number of waypoints during this shrinking, a curve of Pc vs covariance size can be produced; and it will show a maximum Pc value at some point in that sequence. If the aspect ratio of the joint covariance is known or can be reasonably estimated, it is extremely straightforward to find this maximum Pc value through a brute-force technique of shrinking and expanding (usually the) secondary covariance in the CDM.

What does one do with this maximum Pc value once calculated? The recommended operational theory and application of the technique is the following. If the event of interest is presently in the dilution region (which can be determined by calculating the Mahalanobis distance), then the Pc value could potentially be understated due to poor data quality. One can therefore calculate the maximum Pc according to this technique and compare it to the Pc threshold that is typically used to require mitigation actions. If the maximum Pc is below this threshold, then one can conclude that, at least based on the data at hand, the conjunction cannot in fact be dangerous and thus can be safely set aside (at least until the next update, since subsequent updates can change the miss distance and thus the entire maximum Pc curve). If the maximum Pc is above the mitigation threshold, then it is at least conceivable that the conjunction could be dangerous; and a prophylactic mitigation action may be justified merely on this suspicion.

Figure 5 gives maximum Pc results from two CARA historical conjunction events, showing the full Pc possibility curve over a large number of scale factors applied to the secondary object's covariance. The event in the leftmost graph has a nominal Pc of 6.3E-09 and a maximum Pc of 4.6E-06; because this latter value is below the usual mitigation threshold of 1E-04, the event can be dismissed. The rightmost graph shows an event with a nominal Pc of 6.9E-07 and a maximum Pc of 4.6E-04, which is above the mitigation threshold; one must therefore debate how such a case is to be handled.

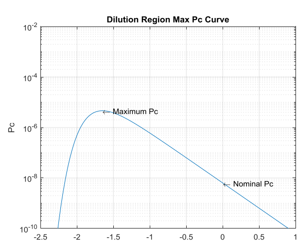

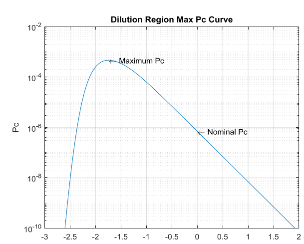

Figure 5: Maximum Pc curves for two different CARA conjunctions

The Maximum Pc parameter and its recommended use argue for a notably different fundamental question and null hypothesis from that assigned to the vanilla Pc calculation. The test works naturally to demonstrate certain situations to be low-risk while remaining ambiguous on those that, with the maximum Pc construct, exceed the mitigation threshold; this sets up a robust situation for a null hypothesis that favors remediation and that can be conclusively rejected based on this test. One would have thus the following null hypothesis: the miss distance is less than the HBR, which would imply that when the maximum Pc is above the remediation threshold, a remediation action should be pursued. A fundamental question that would align with the nature of the test and lead one naturally to the above null hypothesis is of the form: Given the data and assumptions regarding possible values of the covariance, does the maximum Pc value justify dismissal of the event?

"Overlapping" Ellipsoids

While the idea of ensuring adequate separation of space objects by imposing a minimum level of covariance ellipsoid "overlap" has existed for some time, a recent and rigorous development of the concept has been offered by Balch *et al*.[2] The motivating context for the method is the dilution region problem: because a low Pc does not guarantee safety, what calculation should be performed to determine whether a truly safe situations exists? The approach is straightforward: to ensure that the separation of the primary and secondary covariance ellipsoids at a given confidence level be below the chosen HBR; this will ensure that the possibility of collision be kept to a sufficiently small level so as to be discountable. If the ellipsoid overlap is too large, a maneuver is counselled in order to increase the miss distance between the objects to reduce this quasi-overlap below the chosen threshold. All that is required for this calculation are the data contained on the CDM.

The use of this approach shifts the focus nearly entirely from probability to possibility. In the usual Pc calculation, uncertain data result in less certainty of a dangerous situation and thus contribute to the potential dismissing of the event; in this construct, less certainty results in a larger covariance that creates an even greater likelihood of nontrivial quasi-overlap with that for a proximate object and is therefore more likely to counsel a mitigation action. So the mere possibility of a collision is channeled into a recommendation for remediation.

The fundamental question that encapsulates this dynamic, which appears to flow directly from the nature of the construct itself, can be as simply expressed as: Do the data rule out the possibility of a collision? The associated null hypothesis would then take the form: The covariance ellipses encroach on each other to a non-discountable degree. If this null hypothesis cannot be rejected, a mitigation action would be warranted.

Categorization Summary and Conclusions

The salient points of the categorization activity performed above are summarized in Table 1:

Table 1: Categorization Summary of Multiple Risk Assessment Techniques

	Technique	Fundamental Question	Null Hypothesis	Data Required
Probability ↑	Vanilla Pc Calculation	Do the presented data justify a decision to mitigate the conjunction?	The actual miss distance is greater than the hard-body radius	Immediate CDM
	Wald Sequential Probability Ratio Test	*Do the presented data and background risk analysis justify a decision to mitigate the conjunction?*	*The actual miss distance is greater than the hard-body radius*	Immediate CDM + background risk between primary and secondary
	CARA Pc Uncertainty	Given the current data and historical covariance realism information, does the Pc range of values justify a decision to mitigate?	The actual miss distance is greater than the hard-body radius	Immediate CDM + large historical archive of by-object covariance realism data
	Pc Sensitivity	*Given the current data and covariance realism assumptions, do the Pc range of values justify a decision to mitigate?*	*The actual miss distance is greater than the hard-body radius*	Immediate CDM + scale-factor end-points for primary and secondary covariance
	Maximum Pc	Given the data and assumptions regarding possible values of the covariance, does the maximum Pc value justify dismissal of the event?	The actual miss distance is less than the hard-body radius	Immediate CDM + expected covariance aspect ratio
↓ Possibility	Ellipse Overlap	Do the data rule out the possibility of a collision?	The covariance ellipses overlap to a non-discountable degree	Immediate CDM

In examining the results of this categorization in summary, the following items of interest are observed:

- The different major risk analysis techniques and parameters broadly, and reasonably evenly, span the spectrum of probability to possibility. Of the techniques surveyed, two fall strongly in the probability camp, one mostly in the probability camp, two reasonably in the plausibility camp, and one in the possibility camp.
- Some techniques seem intrinsically to favor a particular fundamental question and null hypothesis; others, shown in the table in italics, are more fungible. Because it inherently includes alarm, dismissal, and “no decision” categories, the WSPRT does not really need a fundamental question at all; and if one wishes to require one, then the construct could be pushed in either direction. The Pc Sensitivity technique seems to be flexible as well; as presently used it embraces a null hypothesis of refraining from mitigation, but by construction the approach is more open to mitigation.

- In the midst of the "plausibility" category, the null hypothesis flips from refraining from mitigation to favoring mitigation. Such a development makes sense given the shift of locus from probability to possibility and is probably to some degree a result of the categorization schema.
- The techniques at the extremes of the spectrum require the least amount of conjunction information to drive them. Seeking "plausibility" usually requires looking beyond the immediate situation to identify and evaluate other outcomes, and thus additional data are needed.

The exercise of applying categories to the major risk assessment proposals, especially through their motivation by and different postures towards the dilution region phenomenon, directs needed attention to the question of how risk assessment should respond to situations that do not allow a definitive statement of either high or low risk. Depending on one's definition of "high" and "low", and depending further on anticipated changes to the conjunction assessment environment (such as increased catalogue size due to new sensors, smaller and harder-to-track primaries, large constellations using non-Keplerian orbits, and autonomous conjunction assessment), the ambiguous middle ground is likely to be encountered more and more frequently; and CA practitioners will require a well-constructed framework to respond to such situations. It is primarily for this purpose that a comprehensive CA philosophy needs to be developed that will allow: individual CA groups to adjudicate the competing claims of requirements for the sustainable use of space; operators to protect of individual spacecraft with minimal mission interruption; and cooperative CA activities between operators. In advance of the development of such a philosophy, the categorization schema here developed can be used to help to discuss profitably and negotiate the different strains of CA risk analysis parameters, as well as forming the building blocks and helping to isolate the key concepts needed for this philosophical formulation.

REFERENCES

[1] Carpenter, J.R. *et al.* "Relevance of the American Statistical Society's Warning on p-Values for Conjunction Assessment." 2017 AAS/AIAA Astrodynamics Specialist Conference (Paper # 17-614), Stevenson WA, August 2017.

[2] Balch, M.S., Martin, R., and Ferson, S. "Satellite Conjunction Assessment and the False Confidence Theorem." Alexandria Validation Consulting internal report (arXiv:1706.08565v3 [math.ST] 22 MAR 2018).

[3] Wasserstein, R.L. and Lazar, N.A. "The ASA's Statement on p-Values: Context, Process, and Purpose." *The American Statistician*, Vol. 70, No. 2, 2016, pp. 129–133.

[4] 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).

[5] Patera, R.P. "General Method for Calculation Satellite Collision Probability." *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 4 (July-August 2001), pp. 716-722

[6] Alfano, S. "A Numerical Implementation of Spherical Object Collision Probability." *Journal of the Astronautical Sciences*, Vol. 53 No. 1 (January-March 2005), pp. 103-109.

[7] Chan, F.C. *Spacecraft Collision Probability*. El Segundo, CA: The Aerospace Press, 2008.

[8] Coppola, V.T. "Including Velocity Uncertainty in the Probability of Collision between Space Objects." *AAS/AIAA Spaceflight Mechanics Meeting*, Charleston SC, Paper 12-247, Feb. 2012.

[9] Alfano, S. "Satellite Conjunction Monte Carlo Analysis." *Advances in the Astronautical Sciences* Vol. 134 (2009), pp. 2007-2024.

[10] Hall, D.T., Casali, S.J., Johnson, L.C., Skrehart, B.B., and Baars, L.G. "High-Fidelity Collision Probabilities Estimated using Brute Force Monte Carlo Simulations." 2018 AAS/AIAA Astrodynamics Specialist Conference (Paper #18-244), Snowbird UT, August 2018.

[11] Alfano, S. "Relating Position Uncertainty to Maximum Conjunction Probability." *Journal of the Astronautical Sciences*, Vol. 53 No. 2 (April-June 2005), pp. 193-205.

[12] Hejduk, M.D., Snow, D.E., and Newman, L.K. "Satellite Conjunction Assessment Risk Analysis for "Dilution Region" Events: Issues and Operational Approaches." Fifth Annual Space Traffic Management Conference "Progress through Collaboration, Austin TX, February 2019.

[13] Carpenter, J.R. and Markey, F.L. "Wald Sequential Probability Ratio Test for Space Object Conjunction Assessment." *Journal of Guidance, Control, and Dynamics* Vol. 37 No. 5 (September-October 2014), pp. 1385-1396.

[14] Schilling, B., Taleb, Y., Carptenter, J., Balducci, M., and Williams, T. "Operational Experience with the Wald Sequential Probability Ratio Test for Conjunction Assessment from the Magnetospheric MultiScale Mission." 2016 AIAA/AAS Astrodynamics Specialist Conference, Long Beach CA, August 2016.

[15] Hejduk, M.D. "Space Catalogue Accuracy Modeling Simplifications." 2008 AAS Astrodynamics Specialists Conference, Honolulu, HI, August 2008.

[16] Hejduk, M.D., Casali, S.J., Cappellucci, D.A., Ericson, N.L., and Snow, D.E. "A Catalogue-Wide Implementation of General Perturbations Orbit Determination Extrapolated from Higher-Order Theory Solutions." 2013 AAS/AIAA Space Flight Mechanics Meeting, Kauai, HI. February 2013.

[17] Hejduk, M.D. and Johnson, L.C. "Approaches to Evaluating Probability of Collision Uncertainty." 2016 AAS Space Flight Mechanics Meeting (paper # 16-241), Napa CA, February 2016.

[18] Laporte, F. *et al.* "Conjunction Assessment Short Course." 2018 AMOS Technical Conference, Wailea HI, September 2018.