

Satellite Conjunction Assessment Risk Analysis for “Dilution Region” Events: Issues and Operational Approaches

M.D. Hejduk, Astrorum Consulting LLC (on behalf of The Aerospace Corporation);
mdhejduk@astrorum.us

D.E. Snow, Omitron Corporation; daniel.snow@omitron.com

L.K. Newman, lauri.k.newman@nasa.gov

NASA Conjunction Assessment Risk Analysis Program

Abstract

An important activity within Space Traffic Management is the detection and prevention of possible on-orbit collisions between space objects. The principal parameter for assessing collision likelihood is the probability of collision, which is widely accepted among conjunction assessment practitioners; but it possesses a known deficiency in that it can produce a false sense of safety when the orbital position uncertainties for the conjuncting objects are high. The probability of collision is said to be “diluted” in such a situation and to understate the possible risk; certain approaches have been recommended by researchers to provide (largely conservative) risk estimates and remediation methodologies in these cases. The present analysis explores two of the main proposals for quantifying and remediating possible risk in the dilution region and quantifies their operational implications. These implications with regard to imputed additional workload are considerable, especially in anticipating the conjunction event levels expected with the deployment of the USAF Space Fence radar. This effort has been undertaken as part of a larger enterprise that seeks to clarify the philosophical and statistical underpinnings of the conjunction risk assessment process. The analysis presented herein argues that a form of hypothesis testing is implicitly used in conjunction assessment risk analysis, and that there are a number of conceptual and practical reasons for constructing the associated null hypothesis to counsel against a satellite conjunction remediation action. In short, it is concluded that, for the purposes of determining whether a conjunction remediation action should be pursued, dilution-region probabilities of collision should be treated no differently from those produced under other circumstances.

Introduction

Space Traffic Management (STM) is a collection of a broad set of individual disciplines. Intended to manage the use of Earth-proximate space, it includes studies and modeling to understand space debris production, evolution, and mitigation; satellite design considerations to reduce collision vulnerability and improve debris shielding; on-orbit safety operations for ascent, satellite mission performance, and disposal; active debris removal conceptual and operational development; and legal and policy considerations. The NASA Conjunction Assessment Risk Analysis (CARA) program focuses on the operational on-orbit safety portion of STM, creating algorithms and software to minimize the likelihood of collisions between protected satellites and other space objects and using these algorithms and software operationally to protect NASA payloads. The purpose of these activities is not only to prevent loss of mission for the protected

space asset, which has always been a consideration in conjunction assessment operations; but also to protect certain important orbital corridors from debris pollution that could render them largely to entirely unusable, or at the least create debris density levels that could accelerate a chain reaction of debris-producing collisions and render the corridor unusable in due time. To perform conjunction assessment risk analysis, one monitors close approaches between the protected asset and the remainder of the known space catalogue; when close approaches are discovered, they are analyzed in some depth to determine whether a collision is likely and therefore should engender some type of collision mitigation action by the protected object. The focus, therefore, is to develop and improve an evaluation mechanism that renders some sort of probabilistic evaluation of the collision potential of a space object close encounter.

The original method of evaluating collision risk was to examine the predicted miss distance (MD) between the two objects at their time of closest approach (TCA), and MD values that represented “too close” a pass between two objects (one a protected asset) signaled dangerous conjunctions and therefore candidates for a remediation action, such as a propulsive maneuver by the protected asset to change its orbit to reduce the miss at TCA to an “acceptable level.” The difficulty here is to determine what “too close” and “acceptable level” are, and in examining the problem one quickly realizes that the distances that would correspond to these adjectival phrases vary depending on the uncertainties with which each of the two satellites’ positions are determined. If the two satellites’ position uncertainties at TCA are extremely small, then a miss distance that also typically might seem small could be acceptable so long as it be sufficiently larger than the uncertainties. Conversely, a somewhat larger miss distance could be notably larger than many times the satellites’ combined physical size but, due to larger position uncertainties, still represent a potentially risky situation because many possible renditions of the satellites’ actual positions, when including uncertainty, could place them close enough to collide. To perform a durable assessment of the situation, an evaluation method is needed that considers the orbital uncertainties to produce a probabilistic evaluation of collision likelihood.

The Probability of Collision

In support of the Space Shuttle program, researchers as early as 1992 had developed a method of rendering such a probabilistic calculation.¹ Called the two-dimensional probability of collision (P_c), this approach is a quite rapid, (largely) analytical calculation of the likelihood that the MD between the two objects will be less than a specified value. The approach requires that the collision be of short enough duration that rectilinear motion between the two satellites as well as invariant position covariances for them can be presumed during the encounter; this is not a burdensome set of assumptions, as it is true of most actual conjunctions outside of the geosynchronous regime. There are a number of published treatments that step through this technique’s derivation in detail;² what will follow here is a brief prose description. The approach examines the conjunction at TCA, combining the two objects’ position error volumes into a joint covariance and placing it at one end of the relative position vector (by convention the end with the secondary object), and combining the two objects’ sizes and placing a sphere circumscribing this joint size, with a radius called the “hard-body radius” or HBR, at the other end of the relative position vector (the end with the primary object). From these assumptions, one can conclude that if a collision is to occur, it will take place in the plane perpendicular to the relative velocity

vector and that the marginal component of the probability perpendicular to the plane will approach unity and can thus be ignored. Since the combined covariance is not changing over this interval, one can project the entire situation into this plane (called the “conjunction” or “encounter” plane”) and evaluate the collision likelihood in two-dimensional space. If this plane is constructed to place the secondary object at the origin and the relative miss vector along the x-axis, one obtains the following:




Figure 1: Conjunction-plane rendering of 2-D P_c calculation

The P_c is thus the amount of uncertainty probability density that falls within the HBR circle, as it is those situations that reflect an actual miss distance that will fall within this circle. Since the uncertainty volume (area once projected into the conjunction plane) extends to infinity (although with of course ever-diminishing density), there will always be some probability density falling within the HBR circle, although often the density is low enough that the P_c is essentially 0 to machine numerical integration precision. The P_c calculation, which involves the integral of the Gaussian probability density over the HBR circular area, is extremely fast ($< 1/1000^{\text{th}}$ of a second in even a slow language such as MATLAB) and is thus a quite practical solution. In addition to this cleanliness, the power of such a calculation is evident: the probability of a collision is a concept to which decision-makers can relate relatively easily (especially if they are already involved with risk management of some type), and the P_c can be placed against or combined with other types of mission failure data and thus given appropriate context. The conjunction assessment community has therefore embraced this parameter as the standard collision likelihood calculation, with a typically-employed remediation threshold of $1E-04$ —meaning that a conjunction remediation action, such as a propulsive maneuver, is taken when there is a 1 in 10,000 chance or better of a collision.

“Dilution Region” Defined and Explained

While the concept of the P_c certainly seems straightforward enough, once one has worked with it for some time a curious phenomenon emerges regarding an ambiguity of meaning

with certain classes of low P_c values. In examining Figure 1, one can see that, for a given HBR, there will be a particular joint covariance size that will maximize the amount of covariance probability density that falls within that HBR and thus will similarly maximize the calculated P_c . Because such a P_c maximum exists, growing or shrinking the covariance from this value will produce smaller P_c values. To wit: if the joint covariance is extremely small, there is very little probability density falling within the HBR circle, and the P_c will be small; if the covariance is extremely large, then the probably density is spread out over a large area, and again there is relatively little falling within the HBR circle, so in this case the P_c will also be small. Upon reflection, these two paths to a small P_c value make physical sense. For the first possibility, because there is little uncertainty in the primary and secondary states, it is highly likely that the estimated miss distance will be close to the “true” miss distance; and if this nominal estimated miss distance is somewhat larger than the HBR, then it is highly likely that the actual miss distance is truly larger than the HBR, and the resultant P_c is thus low. This small joint covariance is sometimes called the “robust” region of performance because the calculation of a low P_c in this case is a robust result. For the second possibility, because the relative position of the two satellites is so poorly determined, there is a broad range of possibilities for the actual relative miss vector; and because the miss distance cannot be negative, it is relatively unlikely that the actual miss will be smaller than the HBR; so a low P_c value is thus produced here as well—not because the actual collision risk is known to be low but because so little is known about the satellites’ actual positions that it cannot be concluded that it is high. This is sometimes called the “dilution” region of performance because the P_c has been “diluted” to a low level through high uncertainties in the satellites’ state estimates. Since a low but diluted P_c is produced largely by data uncertainty, it cannot be taken as a warrant that the protected satellite in such a conjunction is safe.

An analogy with car locations in a large parking lot, while not exact, is perhaps helpful here. If two cars are parked in a large parking lot, and we have been told that one car is parked close to one side of the lot and the other close to the other side, then we can conclude with a high degree of confidence that they are not parked next to each other. If, however, all we know is that the two cars are each parked somewhere in the lot, then we can also surmise that it is unlikely that they are parked next to each other, but only because in a large parking lot it is extremely unlikely that any two particular cars will happen to be adjacent to each other. In this latter case, the low probability does not mean the data have shown us the two cars are far apart; it is only a statement of the general unlikelihood of such an alignment if all one knows is that two cars happen to be in the same general area—if better information were available about the two cars’ location, we might learn rather easily that the two cars are in fact adjacent.

One can depict this situation graphically, and the best way to do this is to plot, for a particular conjunction, the P_c vs the ratio of covariance size to miss distance as the covariance size is modulated; such a situation is represented in Figure 2. The situations in the robust and dilution regions differ notably. If, for example, we begin with a large covariance (right side of graph) and systematically shrink it, we begin with a relatively small P_c value, which increases modestly as the covariance is shrunk until a peak is reached; and as the covariance continues to be shrunk past the peak, the P_c drops off precipitously. It is of note that this sequence of events aligns roughly with a typical conjunction event’s dynamics: when the event is first discovered (usually seven or so days before TCA), the covariances are large due to the long propagation time to TCA, the event is in the dilution region, and the P_c is low. As time progresses to TCA,

the joint covariance shrinks because the period of propagation is shorter and because in many cases additional tracking has been obtained (thus also shrinking the covariance), but the miss distance stays about the same; so the P_c increases and eventually reaches a peak. Finally, as TCA is approached, the joint covariance shrinks to the point that the P_c drops off substantially (if it were the rare case of an actual detected collision and no remediation were pursued, the P_c would instead continue to increase up to unity). While this general paradigm is straightforward enough, two issues that make it difficult to apply to actual conjunctions are the fact that the nominal miss distance actually changes each time new data are used in a CA screening (typically multiple times per day), creating a less-than-smooth progression, and often non-progression, along the Figure 2 curve; and the maneuver commitment time, which is the point at which a maneuver decision has to be executed, may not fall within the robust region, complicating the risk assessment process.

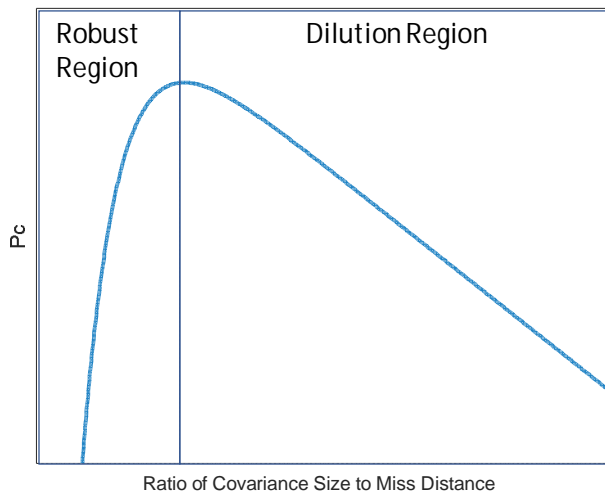


Figure 2: P_c vs the ratio of covariance size to miss distance

This latter comment takes us to the issue that is the focus of this paper: how the CA risk assessment analyst should respond when presented with a P_c value produced in the dilution region, especially if the calculated P_c is low and would typically be seen as certifying a certain degree of safety. Some commentators believe that the robust/dilution region distinction is largely to entirely irrelevant when making remediation decisions³; others believe that the very existence of this problem essentially invalidates the P_c as a conjunction likelihood determination parameter.⁴ The position advanced here is that the utility of the P_c under dilution region conditions is governed by two additional items: 1) the philosophical framework that subtends the CA enterprise, and 2) the viability of alternatives to the P_c that can identify safe conditions yet not interfere excessively with a primary satellite's mission. These two questions will now be explored in depth.

NASA CARA CA Philosophy

Safety is rarely an absolute but rather an incremental attribute; rather than becoming “safe” after a certain amount of fortification, instead there is a graduated increase in safety with the taking of increased precautions. Thus, with CA there is a trade-space between the amount of safety investment (additional funding and additional mission inconvenience and interruption) and the amount of catastrophic risk abatement realized, with no clear threshold indicating when “enough” safety has been procured. In this sense, space safety is rather like automotive safety. Vehicular deaths could be substantially reduced if all known automotive safety features and devices were made mandatory, vehicle weights and reinforcement were substantially increased, and speed limits dramatically lowered; but per-vehicle costs and user inconvenience/invasiveness would be overwhelmingly increased as well. Where one wishes to be on the safety vs cost/inconvenience spectrum is thus a prudential decision that, appropriately, is often more heuristic than technical. In keeping with the automotive safety theme, as an example one can examine the approach used by the National Highway Traffic Safety Administration, the division of the Department of Transportation (DoT) that sets automotive safety policy. When evaluating automotive safety devices, typically years of field data on the performance of the device (at that point vehicle options that some consumers elect to purchase) are examined; and if the device’s safety advantages are shown to be considerable and the cost (and other burdens) relatively manageable and tolerable, DoT will consider making this feature mandatory. This same sort of approach is used by NASA CARA in recommending conjunction assessment safety policies for NASA, the principal difference being that so few satellite collisions occur routinely that policy recommendations must rely more on studies and simulations than field data.

In adjudicating this safety continuum for CA, NASA CARA embraces the following statement of principles: “To take prudent measures, at reasonable cost, to improve safety of flight, without imposing an undue burden on space mission operations.” A number of important points are extractable from this statement:

- Prudence, not analysis or technical argument, is the ultimate governing paradigm in selecting particular safety measures or approaches. The choosing of a risk exposure posture is a human decision that considers as many non-technical as technical factors.
- As there exists a trade-off between safety and cost, there is no mandate to “over-invest” in safety, pushing its cost requirements beyond what is reasonable.
- The burden that the safety activities, typically here meaning excessive orbit readjustments, impose on mission conduct is an extremely important consideration; and while mission accommodations to safety-related activities are certainly expected, a balance must be struck that allows mission activities to continue largely unimpeded.

CA Decisions and Hypothesis Testing

As explained previously, the P_c is calculated from state estimates and covariances for both objects in conjunction, propagated to TCA. Because of this, some individuals see the entire enterprise as merely an estimation problem—one is taking the outputs from an estimation process and making a straightforward calculation, which is then compared to a threshold; there is

no additional statistical scaffolding and complication required. One cannot deny a certain reasonability to this claim, and it without question describes the manner in which many CA practitioners operate: comparing Pc values to thresholds without even considering the activity as part of a statistics problem.

However, there are aspects of the procedure that suggest a statistics context. First (and perhaps somewhat pedantically), the calculation of the Pc itself, while it uses the results of an estimation process, is not itself such a result—it is a subsidiary calculation and therefore not a product of an estimation process *per se*; so it is not correct to assert that the Pc itself is generated by an estimation technique. Second, the comparison of the Pc to a threshold defines a critical region,⁵ leading to certain assumptions and actions; for example, if the Pc exceeds the threshold, then the likelihood that the MD will be smaller than the HBR is too great and a remediation action warranted. Third, while it is common to work with just the Pc and the associated remediation threshold, many practitioners are wishing for (and some attempting to construct) a confidence interval as part of the Pc calculation so one can understand the statistical likelihood of the “true” Pc exceeding a given threshold.⁶ In short, the comparison of the result to a threshold, the accompanying implicit definition of a critical region, and the desire for confidence boundaries transforms the context from merely the observation of a certain calculated result to a statistical framework that suggests the dynamics of a hypothesis test. Let us examine the situation in more depth to determine what a hypothesis test context for CA decisions might look like.

In 2012, V. Coppola offered the following definition of an on-orbit collision: “The miss distance is less than or equal to the hard-body radius.”⁷ In a subsequent publication by J.R. Carpenter *et al.*, this statement was taken up as the null hypothesis for a study effort regarding CA operations,⁸ both because it already existed in print in that form and because it allowed an alignment of form with the usual construction of hypothesis testing: a calculated p-value (for which the Pc is here serving as a proxy) smaller than the threshold (here the remediation threshold) leads to a rejection of the null hypothesis. A small Pc value would therefore allow the rejection of the null hypothesis that the MD is less than the HBR—a statement that can be used accurately to describe CA operational conduct and therefore one that certainly seems reasonable.

While statistics texts often state, correctly, that the researcher has wide latitude in shaping the null hypothesis for a particular problem, they frequently also remark that the null hypothesis is the “typical” or “ordinary” situation,⁹ which one must marshal evidence to displace¹⁰; so the null hypothesis should give the view or course of action that will be embraced in the absence of compelling evidence to reject it. When evaluated in light of these considerations, a null hypothesis that describes a risky situation (“the miss distance is less than the HBR”), for which one would typically pursue conjunction remediation, is strangely formulated, at least in the light of more than a decade of CARA operational experience. There are a number of inherent disadvantages of and risks associated with satellite maneuvers (stuck thrusters, other mechanical failures, gap in mission/science data until nominal orbit restored), so remediation is not seen as the default action but one for which there must be explicit justification or argumentation. Additionally, multiple different aspects of the conjunction dataset must all be in place for a remediation action to be desirable — not only must the Pc be at a worrisome level, but the state estimate data for the particular conjunction must be of a sufficiently high quality to be

considered actionable and the space weather situation between the current time and TCA must be sufficiently stable. These considerations suggest that a more appropriate null hypothesis for CA would be not to remediate; or to follow Coppola's formulary: "The miss distance is greater than the hard-body radius." This articulation more naturally accounts for the inherent risks of remediation actions, and it gives a more natural manner in which remediation is rejected if any of the three criteria listed above is not present. Finally, it also aligns with how seasoned operators view the problem, namely to remediate only if all the aspects of the problem point to a remediation action.

One could, in spite of these arguments, still maintain that the safer course is to remediate whenever any indication at all of a conjunction risk exists, and thus that the null hypothesis should be to presume a collision unless the evidence can disprove it to an acceptable level. To do this, however, would be to ignore the space-debris-infused situation in which satellites are presently placed into service. Current modeling efforts estimate that there are *ca.* 500,000 pieces of space debris 1cm or larger in Earth orbit,¹¹ a size that exceeds the shielding level of all spacecraft, except perhaps those used for human spaceflight, and thus would cause mission loss in a collision. Since the current catalogue size, maintained down to approximately 10 cm, is *ca.* 22,000 objects, about 95% of potential collisions do not even have the possibility of being discovered and remediated; yet we accept this substantial risk and launch satellites anyway. Since the very nature of current satellite operations already accepts that the very great majority of the collision risk cannot be remediated, it does not seem consistent to propose that situations with ambiguous or inconclusive data should default to requiring (often very large) remediation actions rather than simply becoming a small addition to the unremediable collision risk. Instead, situations in which dangerous conjunctions are clearly identified, in the presence of actionable data, are remediated; and other situations are treated in a manner identical to the large number of unknown and thus unremediable events that are sustained every day by dint of simply occupying an Earth orbit. This posture seems to be the approach most consistent with "prudent action, at reasonable cost, to improve safety of flight...."

Given the above, the recommended philosophical approach to CA risk assessment is as follows. The fundamental question that governs the analysis of each event is "Do the presented data provide evidence to justify a decision to remediate?" This framing of the fundamental question places presumption with not remediating: evidence must be presented explicitly to justify a remediation decision, and the absence or questionability of such evidence leaves the situation in the default state of not pursuing a remediation action. The null hypothesis consistent with this fundamental question, as mentioned previously, is "The miss distance is greater than the HBR." To reject this null hypothesis it is necessary to establish that the P_c is greater than the remediation threshold (to an appropriate confidence level, if it is possible to assess this), that the state and covariance data subtending the P_c calculation are durable and thus will produce an actionable P_c evaluation, and the propagation situation is expected to be stable to the point that the propagated states represent a credible estimate of the expected situation at TCA (typically this means a relatively unperturbed space weather environment). If any of these three components to the decision is not present, then there is not sufficient evidence to reject the null hypothesis; thus the justifiable response is to refrain from a remediation action.

Hypothesis Testing and the Dilution Region

At the point at which a conjunction remediation decision must be made, any given event can present a P_c that is above or below a remediation threshold and can also be in the robust or dilution region. Given the above null hypothesis, and presuming for the moment that the state data are actionable and the space weather situation stable, the following matrix would describe the possible remediation decision outcomes:

Table 1: Remediation decision outcomes

	$P_c > \text{Threshold}$	$P_c < \text{Threshold}$
Robust Region	Risky	Safe
Dilution Region	Risky	No Conclusion

When in the robust region, the evaluation is straightforward: because the covariances are relatively small, then most of the miss distance probability density will be close to the estimate of the mean MD; so if the P_c is above the remediation threshold, the situation is seen as risky, and if below the threshold, it is seen as safe. In the dilution region, because the covariances are large due to position uncertainty, much of the miss distance probability density is away from the mean MD, allowing a much broader set of MD values and therefore driving the P_c to a lower value. If despite this “diluting” effect the P_c is still above the threshold, then the situation truly is serious and worthy of remediation (although the remediation action to create a post-action safe situation might need to be larger than one would like, given the sizes of the uncertainties). If the P_c is below the threshold, however, then no durable conclusion regarding safety can be drawn—it might be the case that the two objects will not in fact pass all that closely to each other; or it may be that they actually will pass dangerously close to each other but, because the uncertainties in the data are so great, the reported P_c is low and thus deceptively reassuring.

In this latter case, while it would be possible to remediate such situations preventively, the manner in which the fundamental question and null hypothesis have been formulated establish that there is no requirement to pursue a remediation action in such circumstances. The presence of a low P_c value in the dilution region is testimony to a poor understanding of the orbital situation; and given that the data do not allow a clear conclusion about the situation, one cannot definitively reject the null hypothesis and mandate a remediation action. The situation instead becomes part of the very large holding of collision risk from objects below the tracking threshold and thus unremediable.

The matter is thrown into relief even more strongly by the imminent deployment of the USAF Space Fence radar, a tracking radar with the advertised ability to track objects down to 5cm. It has been hypothesized that, when this new radar is enabled in the latter part of 2019, the space catalogue will increase in size by a factor of anywhere from three to ten; and while there is some variation due to orbital regime, the number of predicted conjunction events is expected to increase by similar levels. Since nearly all of these new objects will be tracked by only a single sensor and thus can be expected to have larger covariances, it is quite likely that many of the CA events caused by these new objects will be in the dilution region. What posture should CA risk assessment take towards such events? To maintain that there is a mandate to remediate such

dilution-region events seems ironic, since before this new radar was enabled these same events were undetectable and therefore sustained each day with the risk unreservedly accepted. The Space Fence did not change the number of objects in space or the expected long-term collision risk for any protected asset; all it did was make operators aware of a larger number of such objects and events. Given this situation, why would there be a mandate to pursue remediation in situations in which new Space Fence data are ambiguous? Instead, one should act to remediate when the evidence is clear and convincing, treating ambiguous cases in a manner similar to the conjunctions presented by the large amount of untracked space debris.

Proposed Risk Evaluation Approaches for the Dilution Region

As argued above, given the fundamental imperative guiding NASA CA (“To take prudent measures, at reasonable cost, to improve safety of flight, without imposing an undue burden on space mission operations”), there is no mandate to remediate conjunctions in the dilution region whose P_c values fall below the usual remediation threshold. At the same time, this imperative does not prohibit the remediation of risk in such cases, so long as such remediation actions not run afoul of the imperative’s final exhortation not to impose an undue burden on space mission operations. In order to formulate a more robust set of best practices for CA, it is necessary to investigate known reasonable proposals for risk remediation in the dilution region to determine whether they could be so employed without imposing operational burdens. While such investigation should be directed at any and all such proposals as they arise, two of the most prominent will be described and examined here.

The *Maximum P_c* approach was developed and presented by S. Alfano in a paper in 2005.¹² Recognizing that by definition events in the dilution region produce an undersized (“diluted”) P_c due to what could be called an oversized covariance, the proposed technique attempts to determine what the largest possible P_c could be, subject to certain constraints, for the event if additional tracking data had been available. It makes the first-order presumption that the effect of these additional tracks would be to shrink the covariance’s overall size while preserving both its aspect ratio and the event’s nominal miss distance. It thus successively shrinks the event’s joint covariance (or merely the secondary covariance, if it is believed that the protected satellite’s covariance is well determined) iteratively, recomputing the P_c with each successive shrinking, until a maximum P_c value is reached. If this maximum P_c is below the remediation threshold, then one can conclude with some confidence that the event is not dangerous (at least until the next state estimate update, at which point the entire situation could be altered). If the maximum P_c is above the remediation threshold, then it is possible that the event could actually be dangerous; one could plan and execute a remediation action to bring this maximum P_c value down to a level considered safe. Depending of the size of this maximum P_c , the needed remediation action might be quite large.




Figure 3: Two examples of Maximum Pc computation for dilution-region events

Figure 3 above shows the possible Pc curves generated by shrinking the nominal joint covariance until a maximum Pc value is reached (actually, here the joint covariance is both increased and shrunk in order to form the larger continuous curve shown); the two plots present two different archived CARA events. The graph on the left shows a situation in which both the nominal Pc ($6.3\text{E-}09$) and the maximum Pc ($4.6\text{E-}06$) are below the typical remediation threshold of $1\text{E-}04$, so the event can be considered safe. The graph on the right represents an event in which the nominal Pc is well below the remediation threshold ($6.9\text{E-}07$) but the maximum Pc is above it ($4.6\text{E-}04$); this is a situation in which one could consider a remediation action based on what one might call the “plausibility” of conjunction that the Maximum Pc construct reveals here. Of course, it should be pointed out that operational experience is not kind to the assumptions that subtend this method: typically, increased tracking both changes the nominal miss distance and results in at least some reshaping of the joint covariance; so it is not clear that the Maximum Pc really does represent an expected maximum, given that these foundational assumptions are likely to be violated with each successive updates.

The *Ellipse Overlap* approach is a title that can describe a number of related proposals, some of which have been raised and attempted over the years but a group of which has been proposed and given a more rigorous theoretical underpinning by M. Balch.¹³ The general idea is to enforce a certain separation of the two objects by minimizing to a stated level the overlap of their covariance ellipsoids. A simplified approach to this that is somewhat easier to visualize because it can be rendered in the conjunction plane is to enforce a separation between the joint covariance and the HBR circle. At a naïve level, what would make one feel safe — to have the HBR circle lie entirely “outside” the joint covariance error volume? In such a case, a collision would not be possible. Of course, covariance uncertainty volumes do not in fact have defined boundaries but extend to infinity; when rendered as an ellipse, what is reflected is the size of a particular confidence region. So what is actually desired is to arrange for a situation in which the joint covariance for a particular confidence interval—90%, 95%, 99%—does not impinge on the HBR circle. The notional plot in Figure 4 represents such a case: the red ellipse is the joint covariance in the conjunction plane (with center at the origin) at a nominal size, such as 1- or 2-

sigma; the amber circle is the hard-body radius circle one miss distance from the center of the joint covariance, and the blue ellipse is a higher-confidence-level rendering of the joint covariance, let us say for the present the 90% confidence ellipse. If the 90% confidence level is a level that a risk assessment analyst judged to be an adequate level of safety, then this situation would represent a safe situation—the joint covariance, at the desired confidence level, just touches the HBR circle. If the HBR circle were inside the blue ellipse, then the judgment would be that there was not adequate separation between the covariance and HBR circle and that a remediation action to increase this separation would be warranted.

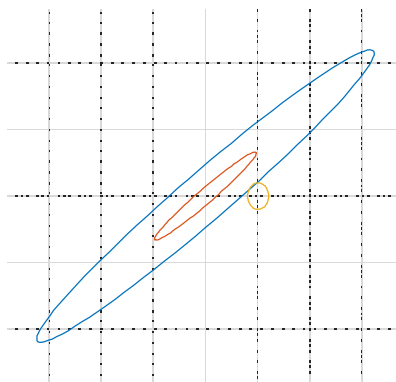


Figure 4: Notional conjunction-plane plot of HBR circle, 1- σ covariance, and 90% confidence covariance

Operational Implications of Dilution Region Evaluation Methods

The two approaches to dilution region risk evaluation described above are both easy to understand conceptually and straightforward to implement; the remaining question is that of the operational implications of the added conservatism of employing either approach. To get a sense of this, both methods were employed against a subset of the 2017 and 2018 CARA historical conjunction database. Conjunctions for twelve protected NASA spacecraft in the near-circular 700 km orbital regime were selected for this profiling exercise. About 11,000 events were examined, originally captured by using a geometric screening volume 0.5 km (radial) x 17 km (in-track) x 20 km (cross-track) about each primary satellite, looking seven days forward into the future. Typically in CARA operations, each event is discovered initially seven days from TCA and then updated three times per day with fresh state estimates and covariances, each propagated to TCA. As explained previously, for most events the risk evolves temporally, beginning low and then increasing to a peak before falling off sharply; so one must choose a particular time-to-TCA at which to exercise the risk evaluation methods. Typically, the remediation action commitment point or the point at which a remediation action decision must be made, is in the neighborhood of 2 days to TCA; for this evaluation, results are tabulated for the 1, 2, and 3 days to TCA points for each event. The great majority of CA events captured by a volumetric screening process end up having a P_c , even a Maximum P_c , of 0 to machine precision throughout the entire 7-day period; so the analysis here focused on dilution-region events with finite P_c values.

Table 2 below gives the comparative results. Each row in the table represents a different remediation action commit point (1, 2, or 3 days to TCA). Column A gives the number of events that would require remediation based on having a “vanilla” Pc value > 1E-04. Column B gives the number of events meriting remediation using the Maximum Pc construct, here with this maximum Pc value exceeding the regular remediation threshold of 1E-04. Column C gives the number of remediation-required events found by enforcing a no-overlap condition between the 95th percentile joint covariance and the HBR circle, both projected into the conjunction plane. Column D gives the ratio of remediation events required using the Maximum Pc construct to the original Pc, and Column E gives the same ratio but instead placing the ellipse overlap results in the numerator.

Table 2: Results of event profiling with different risk assessment approaches

	# of Events Meeting Remediation Criteria				
	A Original Pc	B Maximum Pc	C Ellipse Overlap	D B / A	E C / A
1 Day to TCA	33	75	335	2.3	10.2
2 Days to TCA	48	95	364	2.0	7.6
3 Days to TCA	59	111	378	1.9	6.4

There are several conclusions that can be drawn from these data. First, and most mundane, is that the number of remediation events decreases notably across all three methods as one moves from event consideration at 3 days to TCA to that of 1 day to TCA. This is in conformity with the behavior of the curve that defines the dilution region: that as TCA is approached temporally risk builds to a peak and then falls off precipitously. Second, one observes that the Maximum Pc method, if employed, would approximately double the number of remediation actions for dilution region events, with the ratio becoming slightly more extreme as the remediation action commit time is moved closer to TCA. A doubling of the number of such actions would be felt as a large increase by a number of owners/operators (O/O), although the overall effect on O/O workload is somewhat smaller given that only ~40-45% of the total number of events requiring remediation occur in the dilution region (resulting in an overall multiplicative increase of ~1.6). It may be more telling, however, to focus on the increase likely to be sustained in the Space Fence era, in which the overall catalogue (and therefore number of remediation-requiring events) could increase by a factor of three to ten, and dilution-region events are likely to occupy a larger portion of the entire event set; a doubling of the number of remediation-demanding events, after a previous increase of even a factor of three due to catalogue growth, would be untenable for the mission performance of many satellites, certainly for most of the scientific payloads that NASA operates. So while the event increase due to deploying the Maximum Pc approach could be sustainable at present, it would constitute a poor precedent moving into the Space Fence era.

The ellipse overlap approach produces a much higher increase in remediation-required events, and it is actually relatively insensitive to the confidence interval chosen (90% vs 95% vs 99%). One could bring down the mandatory remediation levels through a substantial lowering of this confidence interval, but it is difficult to accept that a meaningful level of risk abatement has

been achieved if one is satisfied with only a 50% percentile (or lower) confidence ellipse for avoiding HBR circle overlap. Balch himself states that this particular paradigm may not be workable for actual operational implementation, although he expresses hope that researchers can use it as a foundational effort for developing something related but perhaps more serviceable.

Conclusion

The “dilution region” concept and concern has been part of the CA research conversation since its introducing publication, and to date no single set of best practices has emerged to govern how to address conjunction events in this region. By recognizing that the standard method of performing conjunction risk assessment embraces hypothesis-testing mechanics, and establishing that there are strong arguments for formulating the associated null hypothesis as one that counsels against conjunction remediation, one can assemble a compelling case for interpreting low P_c values arising from the dilution region in the same manner as any other P_c , without special consideration for the fact that such a calculation could potentially be overstating the safety of the event. Such arguments only gain strength when considered in the context of the large amount of untrackable debris in which satellites presently operate and for which collision risk is simply accepted—low- P_c , dilution-region events comprise one very small part of the large number of conjunctions with untracked objects that are sustained daily, without knowledge or action. Of the two principal approaches to evaluate and remediate possible collision risk for dilution-region conjunctions, Maximum P_c methods might be deployable with current event densities but could not be sustained with the expected increases that the Space Fence radar will bring; and ellipse overlap methods are likely to be conservative beyond what operations can sustain, both before and after Space Fence deployment. As such, it is CARA’s present recommended practice to treat the P_c , regardless of dilution region positioning, as a durable assessment of collision likelihood for the purposes of considering and selecting remediation actions.

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 See, for example, Chan, F.K. *Spacecraft Collision Probability*, El Segundo, CA: The Aerospace Press, 2007.
 - ³ Frisbee, J.H. “Re-examining Probability Dilution.” 2009 AAS Astrodynamics Specialist Conference (Paper #09-413), Pittsburg PA, August 2009.
 - ⁴ Balch, M.S. “A Corrector for Probability Dilution in Satellite Conjunction Analysis.” 18th AIAA Non-Deterministic Approaches Conference (paper #AIAA 2016-1445), San Diego CA, January 2016.
 - 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).
 - ⁵ DeGroot, M. and Schervish, M. *Probability and Statistics*, 4th edition. Pearson, 2014.
 - ⁶ 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.
 - ⁷ 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.
 - ⁸ 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, Stevenson WA, August 2017.
 - ⁹ Hogg, R.V., Tanis, E.A., and Zimmerman, D.L. *Probability and Statistical Inference*, 9th edition. Pearson, 2015.

¹⁰ Wolfe, D.A. and Schneider, G. *Statistical inference: Estimating Probabilities and Testing and Confirming Models*. Springer, 2014.

Carlton, M.A. and Devore, J.L. *The Basics of Statistical Inference*. Springer, 2017.

¹¹ Liou, J.C. "Overview of the Orbital Debris Problem. 2015 Small Satellite Conference, Logan Utah, August 2015.

¹² Alfano, S. "Relating Position Uncertainty to Maximum Conjunction Probability." *Journal of the Astronautical Sciences*, Vol. 53 No. 2 (April-June 2005), pp. 193-205.

¹³ Balch, *operibus citatis*.