Wald Sequential Probability Ratio Test for Space Object Conjunction Assessment

J. Russell Carpenter[*] and F. Landis Markley[†]

NASA Goddard Space Flight Center, Greenbelt, MD 20771

This paper shows how satellite owner/operators may use sequential estimates of collision probability, along with a prior assessment of the base risk of collision, in a compound hypothesis ratio test to inform decisions concerning collision risk mitigation maneuvers. The compound hypothesis test reduces to a simple probability ratio test, which appears to be a novel result. The test satisfies tolerances related to targeted false alarm and missed detection rates. This result is independent of the method one uses to compute the probability density that one integrates to compute collision probability. A well-established test case from the literature shows that this test yields acceptable results within the constraints of a typical operational conjunction assessment decision timeline. Another example illustrates the use of the test in a practical conjunction assessment scenario based on operations of the International Space Station.

Nomenclature

$\boldsymbol{P}_*$ (Unknown) covariance of the distribution of the true relative position vector at time of closest approach.

$\bar{P}_{fa}$ Target false alarm rate.

$\bar{P}_{md}$ Target missed detection rate.

$\hat{\boldsymbol{P}}_{*|k}$ Covariance of the error in the current estimate, using all available data up to and including time t_k, of the true relative position vector at time of closest approach.

$\hat{\boldsymbol{P}}_{*|o}$ Prior covariance of the distribution of the true relative position vector at time of closest approach.

---

[*]Aerospace Engineer, Navigation and Mission Design Branch, Code 595. Associate Fellow, AIAA.
[†]Aerospace Engineer, Attitude Control Systems Engineering Branch, Code 591. Fellow, AIAA.

$\hat{\boldsymbol{r}}_{*|k}$ Current estimate, using all available data up to and including time t_k, of the distribution of the true relative position vector at time of closest approach.

$\hat{\boldsymbol{r}}_{*|o}$ Prior estimate for the mean of the distribution of the true relative position vector at time of closest approach.

$\mathbb{U}_I$ The set of space object state trajectories that will approach to within the combined hard body radius of the two objects during an interval of interest. An overbar indicates the set of trajectories that will not enter the combined hard body sphere.

Λ_k Likelihood ratio at time t_k.

$\mathrm{p}_{\mathbf{z}}(\boldsymbol{z})$ Probability density function of a random vector $\mathbf{z}$, evaluated at its realization $\boldsymbol{z}$.

$\tilde{\boldsymbol{P}}_{*|k}$ Prediction at time t_k of relative position error covariance at time of closest approach.

$\tilde{\boldsymbol{r}}_{*|k}$ Prediction at time t_k of relative position vector at time of closest approach.

$\boldsymbol{\mu}_*$ (Unknown) mean of the distribution of the true relative position vector at time of closest approach.

$\boldsymbol{x}_t$ The combined position and velocity state vectors of two space objects of interest, relative to the central body they are both orbiting, at time t.

$\boldsymbol{y}_k$ Vector of observations used in the ratio test at time t_k.

$P_{c|k}$ Instantaneous probability of collision, based on information accumulated inclusive of time t_k.

$P_{c|o}$ Instantaneous probability of collision, based on *a priori* information.

P_{fa} Achieved false alarm rate.

P_{md} Achieved missed detection rate.

$q(\boldsymbol{x})$ A function that annihilates a probability density over a specified portion of its domain, in order to define a truncated distributions. An overbar indicates annihilation over the complementary domain.

I.   Introduction

When a maneuverable spacecraft confronts a potentially unsafe conjunction with another space object, its operators must decide whether to maneuver to mitigate the risk of a collision. Such decisions may not be straightforward, since the operators must balance their confidence in the predictions (and their associated formal uncertainties) that detected the conjunction, the actual likelihood of a collision, any risk inherent in performing the maneuver, interruptions to the mission's ongoing operations, and long-term consequences such as depletion of consumable propellant. In principle, operators could quantify their tolerance for performing a maneuver which was not required in terms of an acceptable rate of *false alarms*, and their tolerance for failing to maneuver when a collision was going to occur as an acceptable rate of *missed detections*. Whether such tolerances are explicitly defined, or only

implicitly considered, operators may seek to inform their maneuver decisions using thresholds on various metrics associated with the conjunction.

The most common metric for assessing the risk associated with a conjunction is the collision probability. For example, Foster[1] describes how the International Space Station (ISS) uses collision probability for debris avoidance, stating that "if the collision probability is greater than 10^{-4} and the conjunction geometry has been stable for three of the last four state vector updates, an avoidance maneuver is performed if it is possible." Foster explains how this threshold was derived by balancing "fractional residual risk," i.e. the risk that a debris strike will occur when the probability is below this threshold, against the expected maneuver rate. Jenkin[2] describes a similar method for robotic spacecraft that uses a single collision probability threshold in a sequential procedure: "[i]f the computed collision probability exceeds a given threshold ...[and]... [i]f more accurate orbital data cannot be obtained, a maneuver is necessary." Reference 2 also discusses how large uncertainties in the debris object catalog can lead to excessive maneuver rates, especially for robotic spacecraft that do not receive the same level of consideration as human spaceflight missions. Patera and Petersen[3] also describe a method based on a single collision probability threshold, and give a procedure for how to perform a risk mitigation maneuver. In any such method that uses a single decision threshold, it is not possible to control both the false alarm and missed detection rates.[4] Rather, for any given fixed number of observations, one may control *either* the rate of errors arising from false alarms, *or* the error rate from missed detections; one can only seek to minimize the rate of the uncontrolled error type. This appears to be the motivating principle implicit in works such References 1, 2, and 3. Notably, such methods seek to minimize the total maneuver rate, without explicit consideration of minimizing only the number of maneuvers that did not actually avoid a collision, i.e. the false alarms. This paper proposes an augmentation to such methods, which is to use a Wald Sequential Probability Ratio Test[4] (WSPRT) to inform the collision avoidance decision process. The WSPRT guides decisions based explicitly on false alarm and missed detection criteria. The present work shows that a WSPRT for conjunction assessment reduces to a simple ratio of collision odds.

Many works have considered the problem of computing collision probability for space object conjunctions. Studies of the most typical case, in which (among other assumptions) the relative velocity at the time of closest approach is sufficiently high that one may view the conjunction as instantaneous, have included Foster and Estes,[5] Akella and Alfriend,[6] Patera,[7] Chan,[8] and Alfano.[9] Alfano[10] provides a summary and comparison of many of these methods. Each of these works assumes that one knows, or can accurately approximate, the mean and covariance of the relative position error at the time of closest approach, and that these two moments are sufficient to adequately characterize the probability density. It

is possible to accurately compute a collision probability in such cases, and computing this probability efficiently has been the focus of such works. More general treatments that attempt to approximate the computation of collision probability over some finite interval have included Patera,[11] Chan,[12] McKinley,[13] and Alfano.[14] Alfano[15] compares these methods to Monte Carlo results and provides several useful benchmark problems. More recent work has focused on relaxing the assumption of Gaussian uncertainty, including DeMars et al.[16] using Gaussian mixtures, Jones et al.[17] using polynomial chaos, and Coppola's work[18,19] to carefully address the proper incorporation of velocity uncertainty. These works also assume that all uncertainty may be associated with the initial conditions that produce the conjunction predictions. In Reference 20, the first author of the present work discussed the difficulties that arise if one supposes that random fluctuation (i.e. process noise) occurs in the state errors as they evolve in time. The state of the art in such cases appears to be the Monte Carlo method, although some progress toward more exact solutions continues, as for example Kumar et al.[21] have reported. In any event, the WSPRT compares fixed hypotheses concerning the risk of conjunction. If the state errors can randomly fluctuate, such hypotheses are no longer fixed, and other methods, such as the Shiryaev SPRT[22] may be applicable, but are beyond the scope of the present contribution. Therefore, the sequel is restricted to the case considered by Coppola in Reference 18, which appears to be the most careful treatment to date of the general problem without process noise. Although Coppola's final result depends on a set of Gaussian assumptions, he states that it can be readily extended using Gaussian mixtures, and subsequently DeMars et al.[23] have done so.

The contribution of the present work is to show how to use sequential estimates of collision probability, such as are operationally available from an orbit determination process, in a simple WSPRT that explicitly accounts for decision-maker's preferences, expressed in terms of three quantities: an acceptable rate of false alarms, an acceptable rate of missed detections, and a prior assessment of the risk of collision that is independent of the orbit determination process. The last of these three, interpreted as the base rate of collision over the ensemble of possible realizations of the encounter, plays a central role in the overall decision procedure. While the base rate is in principle unknowable, a suitable proxy may be a collision rate computed from a background debris flux. While References 1, 2, and 3 have used the background rate in a somewhat similar manner, the role of the background rate as a proxy for the base rate in the context of a sequential decision procedure does not appear to have been fully appreciated by the conjunction assessment community up to the present time. After a summary of the work leading to the present result, the paper derives the WSPRT from Coppola's statement of the collision probability problem, discusses some of the implications of this finding, illustrates its methodology with two examples, and finally offers conclusions as to its broad applicability within the conjunction assessment community.

A related paper[24] illustrates this method's application in detail to the Magnetospheric Multi-Scale mission, which at the time of this submission was due to launch in late 2014.

II. Background

This section first reviews the WSPRT, then summarizes prior work by the authors applying the WSPRT to the conjunction assessment problem.

A. Wald's Sequential Probability Ratio Test

The WSPRT is based on associating a set of observations, $\mathbb{Y}_k = \{\boldsymbol{y}_k, \boldsymbol{y}_{k-1}, \ldots, \boldsymbol{y}_1\}$, with a pair of hypotheses about hidden variables that give rise to the observations. These observations are viewed as realizations of a collection of random variables, $\mathbf{Y}_k = \{\mathbf{y}_k, \mathbf{y}_{k-1}, \ldots, \mathbf{y}_1\}$, which has the joint density $\mathrm{p}_{\mathbf{Y}_k}(\boldsymbol{y}_k, \boldsymbol{y}_{k-1}, \ldots, \boldsymbol{y}_1) = \mathrm{p}_{\mathbf{Y}_k}(\mathbb{Y}_k)$. The WSPRT divides the likelihood of these observations having occurred under an alternative hypothesis, $\mathcal{H}_1$, by the likelihood that they have occurred under a null hypothesis, $\mathcal{H}_0$,

$$\Lambda_k = \frac{\mathrm{p}_{\mathbf{Y}_k}(\boldsymbol{y}_k, \boldsymbol{y}_{k-1}, \ldots, \boldsymbol{y}_1 | \mathcal{H}_1)}{\mathrm{p}_{\mathbf{Y}_k}(\boldsymbol{y}_k, \boldsymbol{y}_{k-1}, \ldots, \boldsymbol{y}_1 | \mathcal{H}_0)} \tag{1}$$

It is a general scientific practice to associate the null hypotheses with the condition that one is seeking to disprove. This paper associates the null hypothesis with the condition that the conjunction is unsafe. In a WSPRT, one compares Λ_k to decision limits A and B such that whenever $B < \Lambda_k < A$ one should, if possible, seek another observation. If $\Lambda_k \leq B$, then one should accept the null hypothesis, and in the present case, one would recommend a collision avoidance maneuver. If $\Lambda_k \geq A$, then one should accept the alternative hypothesis, and hence one would dismiss the conjunction alarm[a]. Wald's explanations for the thresholds A and B are that the alternative hypothesis will be accepted if it is A times more likely than the null, and the null hypothesis will be accepted if it is $1/B$ times more likely than the alternative. Wald shows that given enough observations, such a procedure will terminate with probability one, and that the resulting false alarm probability, P_{fa}, and missed detection probability, P_{md}, satisfy the inequalities

$$\frac{1 - P_{fa}}{P_{md}} \geq A \quad \text{and} \quad \frac{P_{fa}}{1 - P_{md}} \leq B \tag{2}$$

---

[a]In the present case, there may be minimal penalty in waiting until all possible measurements have been collected. If the test is still indeterminate at that time, it may be prudent to maneuver, although this would imply an increased false alarm rate.

B. Determination of Decision Limits

Although there do not appear to be analytic methods available to find A and B, this paper follows standard practice in adopting Wald's suggestion to define the decision limits A and B in terms of a target false alarm probability, $\bar{P}_{fa}$, and missed detection probability, $\bar{P}_{md}$, as

$$A = \frac{1 - \bar{P}_{fa}}{\bar{P}_{md}} \quad \text{and} \quad B = \frac{\bar{P}_{fa}}{1 - \bar{P}_{md}} \tag{3}$$

It must be noted that Eqs. (2) and (3) do *not* guarantee that $P_{fa} \leq \bar{P}_{fa}$ and $P_{md} \leq \bar{P}_{md}$. They do provide the weaker inequalities $P_{fa} + P_{md} \leq \bar{P}_{fa} + \bar{P}_{md}$, $P_{fa} \leq \bar{P}_{fa}/(1 - \bar{P}_{md})$, and $P_{md} \leq \bar{P}_{md}/(1 - \bar{P}_{fa})$, however.[4] These inequalities guarantee that at most one of the probabilities P_{fa} or P_{md} can be greater than its target value, and it cannot be much greater in the usual case that both target values are much less than unity.

   This choice is certainly reasonable in the usual case for which $\bar{P}_{fa} + \bar{P}_{md} < 1$, but if an unusual case for which $\bar{P}_{fa} + \bar{P}_{md} \geq 1$ were to occur, Wald's recommendation fails to provide a sensible decision procedure. An alternative to Eq. (3) is to define the decision limits as

$$A' = \frac{1}{\bar{P}_{md}} \quad \text{and} \quad B' = \bar{P}_{fa} \tag{4}$$

These values are very close to those of Eq. (3) if $\bar{P}_{fa} \ll 1$ and $\bar{P}_{md} \ll 1$. The limits of Eq. (4) possess the quality that the maneuver limit depends only on $\bar{P}_{fa}$ and the dismissal limit depends only on $\bar{P}_{md}$. In particular, if $\bar{P}_{fa} \ll 1$, then $B' \ll 1$, meaning that one would rarely alarm, while if $\bar{P}_{fa} \approx 1$, then $B' \approx 1$, and one alarms rather easily. Similarly, if $\bar{P}_{fa} \ll 1$, then $A' \gg 1$, meaning that one rarely dismisses, while if $\bar{P}_{md} \approx 1$, then $A' \approx 1$, meaning that one dismisses rather easily. Equations (2) and (4) can be seen to provide the inequalities $P_{fa} \leq (1 - P_{md})\bar{P}_{fa} \leq \bar{P}_{fa}$ and $P_{md} \leq (1 - P_{fa})\bar{P}_{md} \leq \bar{P}_{md}$, which are stricter bounds than Wald's suggested limits provide. A disadvantage to using A' and B' rather than Wald's suggested A and B is that since $A' \leq A$ and $B' \geq B$, generally more observations will be required to reach a decision. While the achieved false alarm and missed detection rates, P_{fa} and P_{md}, might be expected to exceed the performance requirements implied by $\bar{P}_{fa}$ and $\bar{P}_{md}$, the decrease in efficiency of the test may not outweigh the gain in performance.

C. Prior Applications to Conjunction Assessment

Reference 25 proposed the use of a WSPRT for conjunction avoidance in the context of short-term encounters. In particular, this work followed Chan's approach[8] in assuming that the probability density could be expressed as a Rician. Its development reduced the WSPRT to two forms, both of which involved determination of the Rician's non-centrality parameter

within constrained domains associated with each of the WSPRT hypotheses. One form, denoted the "frequentist" method, involved maximum likelihood estimation of the parameter via a line search, and the other, denoted the "Bayesian" method, involved a convolution with a prior Rician density for the parameter. To reach these results, the work assumed an independent sequence of predictions of the close approach condition. In a numerical example, the Bayesian method significantly outperformed the frequentist method.

The assumption of independent Rician densities was overcome in Reference 26 which reformulated the WSPRT using innovations from a filter bank consisting of two norm-inequality-constrained epoch-state extended Kalman filters. In that approach one filter models a null hypothesis that the miss distance between two conjuncting spacecraft is inside their combined hard body radius at the predicted time of closest approach, and one is constrained by an alternative complementary hypothesis. The epoch-state filter developed for that method explicitly accounts for any process noise present in the system, so long as the process noise does not enter into the states that define the conjunction; for example, it would allow process noise that affects measurement biases, but not process noise that drives position and velocity directly. Because of its epoch-state formulation however, that method still required potentially inaccurate approximations to mapping probability density forward through time. Reference 27 constructed a WSPRT that does not require prediction of probability densities. Instead, one uses solutions to a set of Lambert problems after each measurement update in a sigma-point transformation, to approximate the boundary of the set of current-state velocities that will result in a collision. This boundary is used to define two inequality-constrained *current*-state filters whose innovations formed the likelihood ratio for the Wald test. This method was found to work well when the sigma-point transformation was sufficiently accurate. Although References 26 and 27 both achieved superior performance to that of Reference 25, the filter bank methods require filtering of the measurement data, which may not always be available to conjunction assessment analysts.

Returning to commonly available predictions associated with spacecraft conjunctions, namely a relative state and relative state error covariance at the time of closest approach, Reference 28 found that the likelihood ratio of the WSPRT can be reduced to an especially simple form, involving a particular computation of the current best estimate of collision probability, and a similarly computed estimate of collision probability that is based on prior assumptions about the likelihood of collision. Reference 28 included a conjecture that its results would hold for any computation of collision probability. The conjecture of Reference 28 will be proved in this paper.

III. Problem Statement

View the true position and velocity vectors of two space objects at some time t as random variables denoted as $\mathbf{r}_t^i$, and $\mathbf{v}_t^i$, whose realizations are $\boldsymbol{r}_t^i$, and $\boldsymbol{v}_t^i$, $i = 1, 2$, respectively. The origin of these vectors is the center of mass of the central body the space objects jointly orbit. Collect the position and velocity vectors of the two space objects into the combined state vector, $\mathbf{x}_t = [\mathbf{r}_t^{1\top}, \mathbf{v}_t^{1\top}, \mathbf{r}_t^{2\top}, \mathbf{v}_t^{2\top}]^\top$, whose realization is $\boldsymbol{x}_t = [\boldsymbol{r}_t^{1\top}, \boldsymbol{v}_t^{1\top}, \boldsymbol{r}_t^{2\top}, \boldsymbol{v}_t^{2\top}]^\top$, and which has a probability density function $\mathrm{p}_{\mathbf{x}_t}(\boldsymbol{x}_t)$. Following the development of Reference 18, this density is the solution of a deterministic Fokker-Planck equation, i.e. there is no diffusion (process noise), and hence probability is conserved in the following sense: $\mathrm{p}_{\mathbf{x}_0}(\boldsymbol{x}_0)\,\mathrm{d}\boldsymbol{x}_0 = \mathrm{p}_{\mathbf{x}_t}(\boldsymbol{x}_t)\,\mathrm{d}\boldsymbol{x}_t$.

The problem is to avoid the condition that, at any time in some interval $\mathbb{T}_* = [t_0, t_0 + T]$, the hard body volumes of the two objects overlap. Often, operators may choose to inflate the combined hard body volume(s), for example, by enclosing each object's actual hard body volume with a sphere, which circumscribes the largest extent of the object. Patera[29] has suggested that volumes much larger than a circumscribing sphere may also be useful. However the hard body volumes are defined, it is taken as given that the two objects are not within this common region during times outside the interval $\mathbb{T}_*$. If the combined hard body volume is taken to be a sphere of radius R, then Coppola[18] shows that the sets $\mathbb{V}_0$ and $\mathbb{U}_I$ contain all the states that will collide for a particular realization of the conjunction associated with the random process $\mathbf{x}(t) = \{\mathbf{x}_t \,\forall\, t \in \mathbb{T}_*\}$, where

$$\mathbb{V}_0 = \{\boldsymbol{x}_0 \mid \|\boldsymbol{r}_0^2 - \boldsymbol{r}_0^1\| \leq R\} \tag{5}$$

and

$$\mathbb{U}_I = \bigcup_{t \in (t_0, t_0+T]} \left\{\boldsymbol{x}_t \mid \|\boldsymbol{r}_t^2 - \boldsymbol{r}_t^1\| = R, \|\boldsymbol{r}_\tau^2 - \boldsymbol{r}_\tau^1\| > R \,\forall\, \tau \in [t_0, t)\right\} \tag{6}$$

In the sequel, t_0 is always chosen to be far enough in advance of the conjunction that the probability of collision at t_0 is negligible, and hence only the set $\mathbb{U}_I$ needs to be considered.

If a collision is likely, one or both space objects can be maneuvered to avoid it. At the same time, unnecessary maneuvers should be avoided. A probability of missed detection quantifies the tolerance for failing to maneuver when a maneuver was needed[b]. A false alarm probability quantifies the tolerance for maneuvering when it was unnecessary[c]. The aim of

---

[b]The probability of missed detection is the conditional probability of dismissing a conjunction when a collision avoidance maneuver was necessary; in frequentist terms, it is the ratio of dismissals to the the total number of unsafe conjunctions.

[c]The probability of false alarm is the conditional probability of recommending a collision avoidance maneuver when the maneuver was unnecessary; in frequentist terms, it is the ratio of alarms to the the total

this paper is to find a decision procedure that will guide maneuver decisions, in the context of meeting the specified missed detection and false alarm rates.

IV. Problem Solution

This paper's solution to the collision avoidance problem employs the WSPRT. In the present case, the set of observations for the WSPRT, $\mathbb{Y}_k$, consists of observation vectors $\boldsymbol{y}_k$, each of which contains one or more tracking data observables, for one or both of the space objects, at each time t_k, where $t_k < t_* \, \forall \, k$, and $t_* \in \mathbb{T}_*$ denotes the time of closest approach.

The WSPRT ratio divides the likelihood of these observations having occurred under the alternative hypothesis, $\mathcal{H}_1$, that the conjunction is safe, by the likelihood that they have occurred under the null hypothesis, $\mathcal{H}_0$, that the conjunction is unsafe:

$$\Lambda_k = \frac{\mathrm{p}_{\mathbf{Y}_k}(\boldsymbol{y}_k, \boldsymbol{y}_{k-1}, \ldots, \boldsymbol{y}_1 | \mathcal{H}_1)}{\mathrm{p}_{\mathbf{Y}_k}(\boldsymbol{y}_k, \boldsymbol{y}_{k-1}, \ldots, \boldsymbol{y}_1 | \mathcal{H}_0)} = \frac{\mathrm{p}_{\mathbf{Y}_k}(\mathbb{Y}_k | \mathbf{x}_t \notin \mathbb{U}_I)}{\mathrm{p}_{\mathbf{Y}_k}(\mathbb{Y}_k | \mathbf{x}_t \in \mathbb{U}_I)} \tag{7}$$

According to Eq. (3.32) in Maybeck,[30] the conditional densities required for the likelihood ratio are given by the following ratios:

$$\mathrm{p}_{\mathbf{Y}_k}(\mathbb{Y}_k | \mathbf{x}_t \notin \mathbb{U}_I) = \frac{\int_{\bar{\mathbb{U}}_I} \mathrm{p}_{\mathbf{Y}_k, \mathbf{x}_t}(\mathbb{Y}_k, \boldsymbol{\xi}_t) \, \mathrm{d}\boldsymbol{\xi}_t}{\int_{\bar{\mathbb{U}}_I} \mathrm{p}_{\mathbf{x}_t}(\boldsymbol{\xi}_t) \, \mathrm{d}\boldsymbol{\xi}_t} \tag{8}$$

$$\mathrm{p}_{\mathbf{Y}_k}(\mathbb{Y}_k | \mathbf{x}_t \in \mathbb{U}_I) = \frac{\int_{\mathbb{U}_I} \mathrm{p}_{\mathbf{Y}_k, \mathbf{x}_t}(\mathbb{Y}_k, \boldsymbol{\xi}_t) \, \mathrm{d}\boldsymbol{\xi}_t}{\int_{\mathbb{U}_I} \mathrm{p}_{\mathbf{x}_t}(\boldsymbol{\xi}_t) \, \mathrm{d}\boldsymbol{\xi}_t} \tag{9}$$

where $\bar{\mathbb{U}}_I$ is the complement of $\mathbb{U}_I$, so that the likelihood ratio is given by

$$\Lambda_k = \frac{\int_{\bar{\mathbb{U}}_I} \mathrm{p}_{\mathbf{Y}_k, \mathbf{x}_t}(\mathbb{Y}_k, \boldsymbol{\xi}_t) \, \mathrm{d}\boldsymbol{\xi}_t}{\int_{\bar{\mathbb{U}}_I} \mathrm{p}_{\mathbf{x}_t}(\boldsymbol{\xi}_t) \, \mathrm{d}\boldsymbol{\xi}_t} \frac{\int_{\mathbb{U}_I} \mathrm{p}_{\mathbf{x}_t}(\boldsymbol{\xi}_t) \, \mathrm{d}\boldsymbol{\xi}_t}{\int_{\mathbb{U}_I} \mathrm{p}_{\mathbf{Y}_k, \mathbf{x}_t}(\mathbb{Y}_k, \boldsymbol{\xi}_t) \, \mathrm{d}\boldsymbol{\xi}_t} \tag{10}$$

Recognizing that

$$\int \mathrm{p}_{\mathbf{Y}_k, \mathbf{x}_t}(\mathbb{Y}_k, \boldsymbol{\xi}_t) \, \mathrm{d}\boldsymbol{\xi}_t = \int \mathrm{p}_{\mathbf{x}_t}(\boldsymbol{\xi}_t | \mathbb{Y}_k) \, \mathrm{p}_{\mathbf{Y}_k}(\mathbb{Y}_k) \, \mathrm{d}\boldsymbol{\xi}_t = \mathrm{p}_{\mathbf{Y}_k}(\mathbb{Y}_k) \int \mathrm{p}_{\mathbf{x}_t}(\boldsymbol{\xi}_t | \mathbb{Y}_k) \, \mathrm{d}\boldsymbol{\xi}_t \tag{11}$$

number of safe conjunctions.

the likelihood ratio becomes

$$\Lambda_k = \frac{\int_{\bar{\mathbb{U}}_I} p_{\mathbf{x}_t}(\boldsymbol{\xi}_t \mid \mathbb{Y}_k)\, d\boldsymbol{\xi}_t}{\int_{\bar{\mathbb{U}}_I} p_{\mathbf{x}_t}(\boldsymbol{\xi}_t)\, d\boldsymbol{\xi}_t} \frac{\int_{\mathbb{U}_I} p_{\mathbf{x}_t}(\boldsymbol{\xi}_t)\, d\boldsymbol{\xi}_t}{\int_{\mathbb{U}_I} p_{\mathbf{x}_t}(\boldsymbol{\xi}_t \mid \mathbb{Y}_k)\, d\boldsymbol{\xi}_t}$$
$$= \frac{1 - \int_{\mathbb{U}_I} p_{\mathbf{x}_t}(\boldsymbol{\xi}_t \mid \mathbb{Y}_k)\, d\boldsymbol{\xi}_t}{\int_{\mathbb{U}_I} p_{\mathbf{x}_t}(\boldsymbol{\xi}_t \mid \mathbb{Y}_k)\, d\boldsymbol{\xi}_t} \frac{\int_{\mathbb{U}_I} p_{\mathbf{x}_t}(\boldsymbol{\xi}_t)\, d\boldsymbol{\xi}_t}{1 - \int_{\mathbb{U}_I} p_{\mathbf{x}_t}(\boldsymbol{\xi}_t)\, d\boldsymbol{\xi}_t} \tag{12}$$

Finally, introducing the notation

$$P_{c|k} = \int_{\mathbb{U}_I} p_{\mathbf{x}_t}(\boldsymbol{\xi}_t \mid \mathbb{Y}_k)\, d\boldsymbol{\xi}_t$$
$$P_{c|o} = \int_{\mathbb{U}_I} p_{\mathbf{x}_t}(\boldsymbol{\xi}_t)\, d\boldsymbol{\xi}_t \tag{13}$$

leads to the following expression for the likelihood ratio:

$$\Lambda_k = \frac{1 - P_{c|k}}{P_{c|k}} \frac{P_{c|o}}{1 - P_{c|o}} \tag{14}$$

Remark. Any Bayesian estimator for the state given the observations has $p_{\mathbf{x}_t}(\boldsymbol{x}_t \mid \mathbb{Y}_k)$ as its probability density function. Therefore, $P_{c|k}$ is the estimated probability of collision based on the observation sequence $\mathbb{Y}_k$. The probability $P_{c|o}$ can be viewed as the true underlying probability of collision; in practice, $P_{c|o}$ would be computed using an *a priori* probability density that is not based on any of the observations, but rather on prior assumptions or beliefs about the risk of collision. Section V.B discusses the determination of $P_{c|o}$ in practice.

Remark. Eq. (14) is the result found in Reference 28, but this derivation does not assume that the observations are independent or Gaussian, hence proving the conjecture of Reference 28 that Eq. (14) gives the WSPRT for any method of computing collision probability under fixed hypotheses. The appendix shows that under the Gaussian and short encounter assumptions, Eq. (14) involves the usual methods of batch or sequential estimation of the state and its error covariance to compute $P_{c|k}$, and that $P_{c|o}$ is a similar computation from the estimator's prior state and covariance.

Remark. It is apparent (since $(1 - x)/x$ is monotone decreasing) that

$$\Lambda_k = 1 \Leftrightarrow P_{c|k} = P_{c|o} \tag{15}$$

The left-hand equality means that the observations, $\mathbb{Y}_k$, are equally likely under either hypothesis $\mathcal{H}_0$ or $\mathcal{H}_1$. The right-hand equality means that the observations do not change the prior estimate of the probability of collision.

Eliminating the likelihood ratio as follows can simplify the decision procedure further.

Solving Eq. (14) for $P_{c|k}$ gives

$$P_{c|k} = \frac{P_{c|o}}{\Lambda_k + (1 - \Lambda_k)P_{c|o}} \tag{16}$$

Since the WSPRT will recommend a maneuver if $\Lambda_k \leq B$ and recommend a dismissal if $\Lambda_k > A$, the decision thresholds as thresholds on $P_{c|k}$ can be written as follows:

$$P_{c|k} \geq \frac{P_{c|o}}{B + (1 - B)P_{c|o}} \Rightarrow \text{Alarm (Maneuver)} \tag{17}$$

$$P_{c|k} < \frac{P_{c|o}}{A + (1 - A)P_{c|o}} \Rightarrow \text{Dismiss} \tag{18}$$

Equation (3) leads to a decision expressed solely in terms of the four probabilities $P_{c|k}$, $P_{c|o}$, $\bar{P}_{md}$, and $\bar{P}_{fa}$:

$$P_{c|k} \geq P_c^A = \frac{(1 - \bar{P}_{md})P_{c|o}}{\bar{P}_{fa} + (1 - \bar{P}_{md} - \bar{P}_{fa})P_{c|o}} \tag{19}$$

$$\Rightarrow \text{Alarm (Maneuver)}$$

$$P_{c|k} < P_c^D = \frac{\bar{P}_{md}P_{c|o}}{1 - [\bar{P}_{fa} + (1 - \bar{P}_{md} - \bar{P}_{fa})P_{c|o}]} \tag{20}$$

$$\Rightarrow \text{Dismiss}$$

Figure 1 shows the range of decision limits for a given $P_{c|o}$. The upper, red-hued surface represents the alarm threshold, and the lower, blue-hued surface the dismissal threshold. At the corner in which $\bar{P}_{md} = \bar{P}_{fa} = .5$, the alarm and dismissal thresholds converge to the value of $P_{c|o}$. Figure 1 shows how reducing $\bar{P}_{md}$ from this corner while holding $\bar{P}_{fa}$ fixed results in fairly flat alarm limit, but a rapidly declining dismissal limit. Correspondingly, reducing $\bar{P}_{fa}$ from this corner while holding $\bar{P}_{md}$ fixed results in fairly flat dismissal limit, but a rapidly increasing alarm limit.

Figure 2 illustrates these decision limits for a selection of $\bar{P}_{md}$ and $\bar{P}_{fa}$ values. One may observe that when $\bar{P}_{md} = \bar{P}_{fa}$, a fairly symmetric pair of decision limits occur. Allowing relatively more false alarms than missed detections, which the upper right subplots show, results in a tighter alarm bound, which would tend to allow an alarm to occur with relatively less evidence. The lower left subplots show the symmetric case when one permits relatively more missed detections. The lower right corner of the figure shows the "corner" case described above for which $\bar{P}_{md} = \bar{P}_{fa} = 50\%$. In fact, in all cases for which $\bar{P}_{md} + \bar{P}_{fa} = 1$, the decision limits converge to the value of $P_{c|o}$; but as previously mentioned, such cases are not

Figure 1. Alarm (upper sheet) and dismissal (lower sheet) thresholds on $P_{c|k}$, as a function of $\bar{P}_{md}$ and $\bar{P}_{fa}$, for a particular $P_{c|o}$.

sensible, since error rates greater than 50% are worse than guesses.

From Figure 2 one may also infer the method's sensitivity to an incorrect value of $P_{c|o}$. Considering the center subplot, one may take the true value as $P_{c|o} = 10^{-2}$. Were one to incorrectly assume that $P_{c|o} = 10^{-1}$, that is, too large, one would tend to be too reluctant to alarm, and too eager to dismiss, relative to the thresholds associated with the true $P_{c|o}$. Similarly, were one to incorrectly assume that $P_{c|o} = 10^{-3}$, that is, too small, one would tend to be too reluctant to dismiss, and too eager to alarm, relative to the correct thresholds. Since in practice it is difficult to precisely estimate $P_{c|o}$, it may be advisable to "hedge" the WSPRT as Figure 3 illustrates. By using an upper confidence limit on $P_{c|o}$ for the alarm limit, and a lower confidence limit for the dismissal limit, one expands the region of indecision, and potentially delays the decision until more data makes the case for alarm or dismissal conclusively. Since in the conjunction assessment setting one would always alarm when the WSPRT is inconclusive and there is no more opportunity to collect additional data, such a hedging strategy can only increase false alarms.

Figure 4 indicates how operators could incorporate the WSPRT into an existing conjunction assessment work flow which uses collision probability as one of its decision parameters. As the diagram shows, the test is minimally intrusive, only requiring some additional planning activity to establish acceptable values for the target false alarm and missed detection rates. The dashed line connecting Boxes 1c and 1d highlights the connection between these targets and the background rate of collision, which the section below further discusses. With the distinct WSPRT alarm and dismissal thresholds in hand, the only difference from current
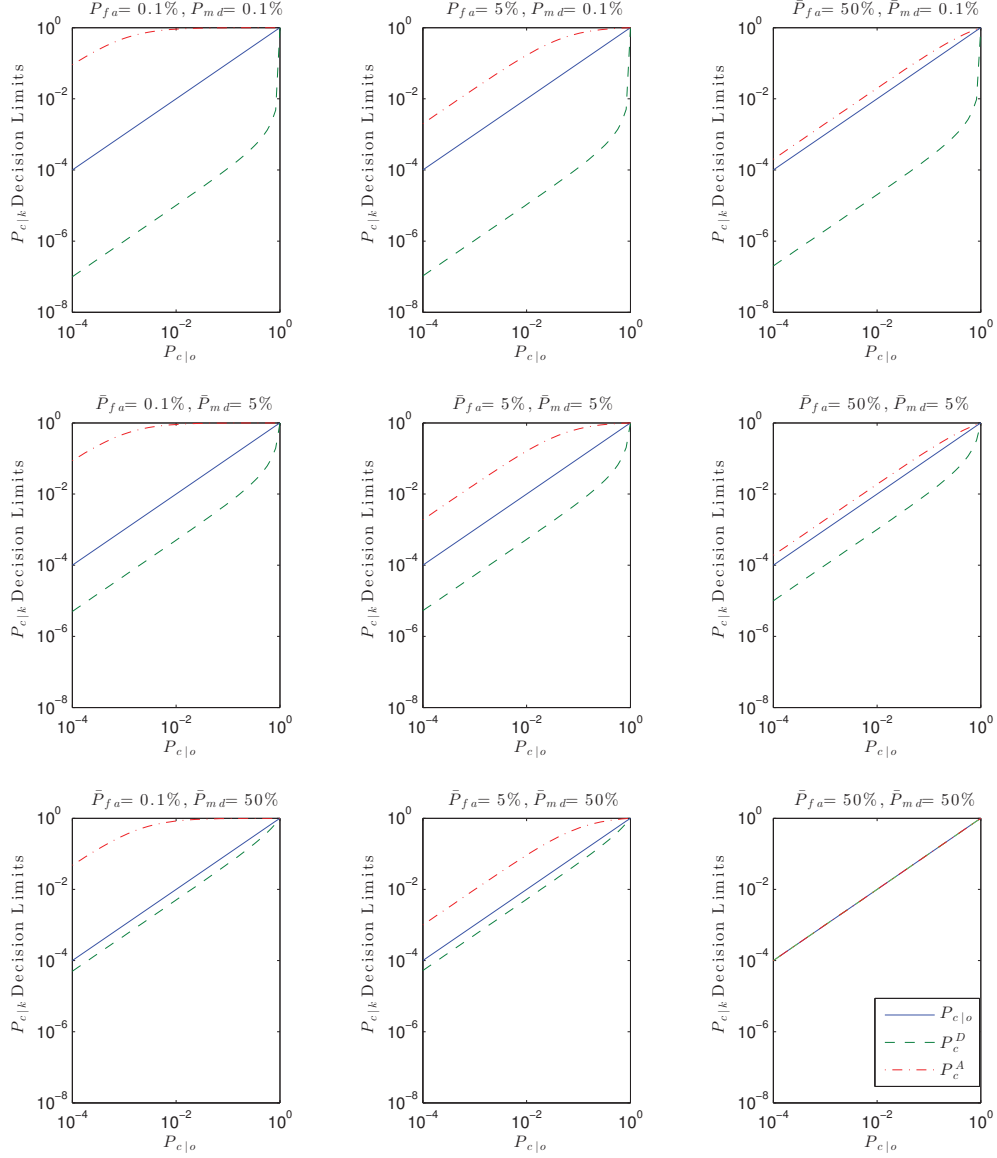
Figure 2. Alarm and dismissal limits on $P_{c|k}$, as a function of $P_{c|o}$. For a given $P_{c|o}$, if $P_{c|k} \geq P_c^A$, the WSPRT suggests a maneuver, and if $P_{c|k} < P_c^D$, the WSPRT suggests dismissing the conjunction.

Figure 3. Hedging to accommodate uncertainty in $P_{c|o}$, indicating how confidence limits on $P_{c|o}$ might be used to hedge decision thresholds.

operations is that the comparisons in Boxes 6a and 6b each use a distinct threshold, rather than sharing a common threshold. Although in principle current operations would have no need to loop back to the orbit determination (OD) stage, in practice operators rarely make a decision based on a single collision probability. Thus, current operations already use an informal sequential procedure, and the WSPRT thresholds merely augment current practice with additional structure.

A possible difference between the work flow Figure 4 illustrates, and some other current conjunction assessment operations, concerns the OD stage. This difference is crucial to the WSPRT. As previously remarked, the derivation above relies on the fact that $p_{\mathbf{x}_t}(\boldsymbol{x}_t \mid \mathbb{Y}_k)$ is produced by any Bayesian estimator utilizing all of the available data, and Figure 4 is consistent with such practice. For instance, the OD process might utilize an extended Kalman filter that processes all of the tracking data sequentially as it arrives. Alternatively, operators could utilize a "sequential-batch" process, which processes each arc of tracking data as a batch, constrained by the propagated covariance from a previous batch. Either approach will result in estimated moments of $p_{\mathbf{x}_t}(\boldsymbol{x}_t \mid \mathbb{Y}_k)$. However, if each round of OD is independent, as for example occurs when each batch estimation cycle uses infinite *a priori* covariance, then the resulting state estimates only reflect information from the tracking data in the particular arc associated with that estimation cycle. In such cases, one may still use the WSPRT, but one must combine the independent estimates in the manner described in the appendix.
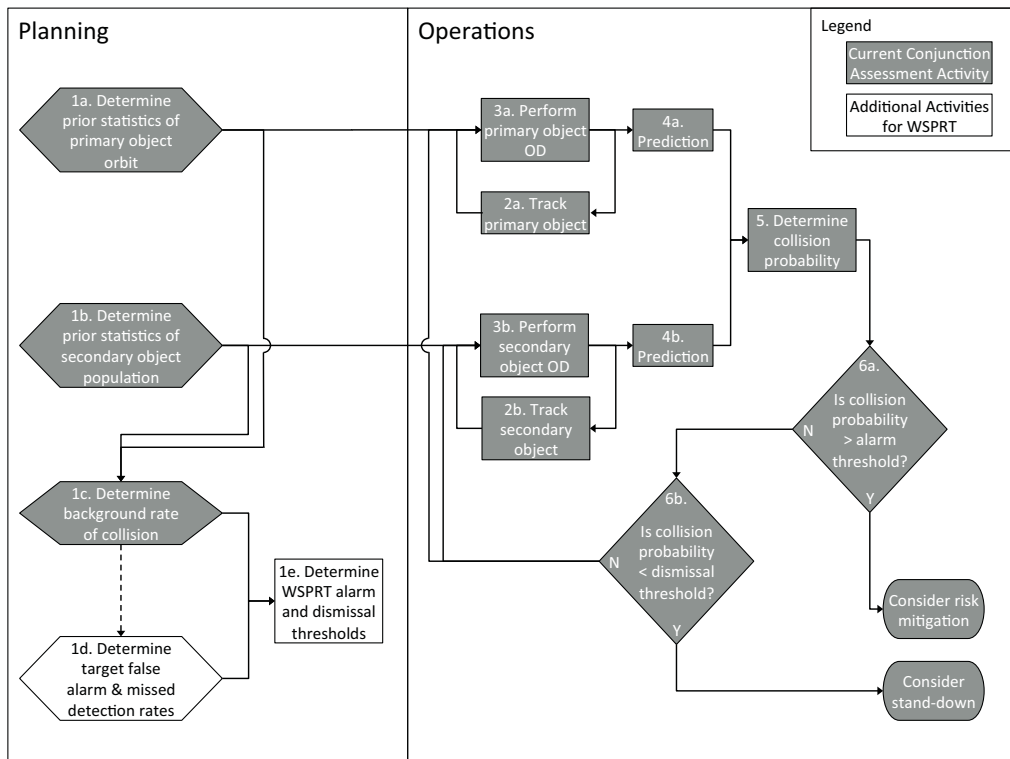
Figure 4. Work flow for use of WSPRT in conjunction assessment operations.

V.   Discussion

The solution described above highlights the central role of the base rate of collision, $P_{c|o}$, in the overall collision avoidance strategy. This section discusses some concerns that may arise when considering extreme values of $P_{c|o}$, and how it might be determined in practice.

A.   Implications of Extreme Base Rates

Possible critiques of the dependence of the decision thresholds on $P_{c|o}$ raise the concern that if $P_{c|o}$ is very near zero, or very near one, thresholds that seem unrealistic occur. However, the thresholds only seem unrealistic because such critiques fail to consider the implications of such extreme base rates. As the following example makes clear, if $P_{c|o}$ is very near zero, the ratio of false alarms to total misses is very different from the ratio of false alarms to total alarms, and the source of the critique is a conflation of these two ratios. Referring to Table 1, consider the case of $P_{c|o} = 1 \times 10^{-6}$: for every million times such a conjunction occurs, one could expect only one close approach inside the combined hard body radius. In such a case, one would expect a decision procedure that allows a false alarm rate of 5% to produce about 50,000 alarms, of which typically only one would not be false, so that the ratio of false alarms to total alarms is nearly 100%. In practice, one would use a screening procedure to eliminate such low probability of collision events from consideration. Similarly, if $P_{c|o}$ is very near one, the ratio of missed detections to total hits is very different from the ratio of missed detections to total dismissals, and once again the source of the critique is a conflation of these two ratios. Referring to Table 2, consider the case of $P_{c|o} = 1 - 1 \times 10^{-6}$: for every million times such a conjunction occurs, one could expect only one close approach outside the combined hard body radius. In such a case, one would expect a decision procedure that allows a missed detection rate of 0.1% to produce about 1,000 dismissals, nearly all of which would result in hits, so that the ratio of missed detections to total dismissals would be nearly 100%. In practice, the mission design process would preclude such high probability of collision events from occurring.

Table 1. Extreme Base Rate Example 1: $P_{c|o}$ close to zero.

	Total Hits = 1	Total Misses = 999,999	
Total Alarms = 50,000	True Alarms = 1	False Alarms = 49,999	$\frac{\text{False Alarms}}{\text{Total Alarms}} = 99.998\%$
Total Dismissals = 950,000	False Dismissals = 0	True Dismissals = 950,000	$\frac{\text{False Dismissals}}{\text{Total Dismissals}} = 0\%$
	$\frac{\text{False Dismissals}}{\text{Total Hits}} = 0\%$	$\frac{\text{False Alarms}}{\text{Total Misses}} = 4.999\%$	

Table 2.   Extreme Base Rate Example 2: $P_{c|o}$ close to one.

	Total Hits = 999,999	Total Misses = 1	
Total Alarms = 999,000	True Alarms = 999,000	False Alarms = 0	$\frac{\text{False Alarms}}{\text{Total Alarms}} = 0\%$
Total Dismissals = 1,000	False Dismissals = 999	True Dismissals = 1	$\frac{\text{False Dismissals}}{\text{Total Dismissals}} = 99.9\%$
	$\frac{\text{False Dismissals}}{\text{Total Hits}} = 0.0999\%$	$\frac{\text{False Alarms}}{\text{Total Misses}} = 0\%$	

The source of confusion in such arguments is a failure to distinguish among two pairs of similar-seeming conditional probabilities. In the false alarm case, these are the probability of an alarm, given that a miss has occurred, $\Pr(\text{alarm}|\text{miss}) \approx \frac{\text{False Alarms}}{\text{Total Misses}}$, and probability of a miss, given that an alarm has occurred, $\Pr(\text{miss}|\text{alarm}) \approx \frac{\text{False Alarms}}{\text{Total Alarms}}$. In the missed detection case, these are the probability of a dismissal, given that a hit has occurred, $\Pr(\text{dismissal}|\text{hit}) \approx \frac{\text{False Dismissals}}{\text{Total Hits}}$, and probability of a hit, given that a dismissal has occurred, $\Pr(\text{hit}|\text{dismissal}) \approx \frac{\text{False Dismissals}}{\text{Total Dismissals}}$. In each case, it is only the former conditional probabilities, which are the false alarm and missed detection rates, that are directly controlled by the decision procedure.

B.   Determination of the Base Rate in Practice

Because in practice the true relative position distribution is of its nature unknowable, the solution found in this paper relies on prior information to establish $P_{c|o}$, which raises the question of the source of such prior knowledge. There are at least two techniques already in use by the conjunction assessment community that might provide a source of such information: flux analysis and screening volumes[d].

McKinley, et al.[32] describe an approach to flux analysis in which

> [t]he debris field is modeled using Two Line Elements (TLEs) from the publicly available General Perturbations catalog ... [which] represent all publicly available data for debris objects and operational spacecraft tracked by [United States Strategic Command].
>
> To determine the steady state debris flux, the ... reference orbit and all the objects in the catalog are propagated, and all objects passing within 50 kilometers of the [reference] space object are recorded. The flux is then calculated as the number of objects penetrating the 50 km sphere per unit of time. The propagation continues until this flux level converges on a steady-state value.

---

[d]Newman[31] provides additional context on NASA's procedures and policies for screening and conjunction assessment for robotic missions.

The 50 km sphere size was chosen based on past experience of the [Goddard Space Flight Center] Conjunction Assessment team performing regime characterization analysis.

The flux ratio method then divides the resulting simulated flux by the flux estimated for a set of one or more operational missions that have actually performed risk mitigation maneuvers to avoid conjunctions. One may then use the flux ratio to scale the actual maneuver rate for the operational mission, as a means of predicting the maneuver rate for the simulated mission. One could then estimate an upper bound for the base rate of collision for the simulated mission by multiplying the predicted maneuver rate by the mission duration. If data were available to characterize the false alarm rate of the operational missions, one could further refine the base rate of collision estimate.

Screening volume analysis, which Narvet, et al.[33] describe, proceeds along somewhat similar lines, but rather than predicting maneuver rates, the goal is to define a volume that will allow detection, using simple and fast methods, of any conjunctions that deserve further scrutiny. As in flux analysis, the method uses a catalog of resident space objects along with a reference trajectory for the mission of interest to determine the components of the miss vectors for any conjunctions that occur in the simulation. These data form empirical cumulative distribution histograms for each component, so that the analyst may determine a threshold for each component that captures the desired proportion of data. In effect, the result is an empirical approximation to the true miss distance distribution. One may use the empirical distribution directly, or fit a distribution of choice to the data, and integrate over the hard body to get the base rate of collision.

Both of these approaches assume that the base rate of collision for all conjunctions for the catalog under study is essentially the same for all conjunctions. For a large catalog, this may be the only practical course of action[e]. Clearly though, there are many cases for which such an assumption fails, e.g. pairs of objects that normally maintain safe separations, but which have a resonance that produces dangerously close approaches at infrequent intervals. In such cases, a determination of the unique $P_{c|o}$ for each conjunction may be desirable. For example, suppose one has a model of the event that precipitated a sequence of close approaches. For specificity, assume the model is a probability density describing the dispersions that may arise from a maneuver by one spacecraft to remain in formation with another spacecraft. Predicting this density to each of the close approaches subsequent to the maneuver and integrating over the combined hard body gives $P_{c|o}$. This method is similar to what would be used to compute $P_{c|k}$ each time an observation became available after the maneuver, but differs in that no post-maneuver information is used to compute $P_{c|o}$ (although the information used to compute $P_{c|o}$ could clearly be used as prior information for the estimates

---

[e]However, one should keep in mind the hedging discussion of the previous section.

that produce the sequence of $P_{c|k}$). In Reference 24, Wawrzyniak, et al. describe how the Magnetospheric Multi-Scale (MMS) mission, due to launch in late 2014, plans to use such an approach.

VI. Examples

Two examples demonstrate the procedure this work advocates. The first example is a reference case adapted from the literature, which has a high collision probability that would nearly always dictate a maneuver. The second example is adapted from operations of the International Space Station, and indicates how the present method could be applied to benefit a current operational scenario. As previously noted, Reference 24 describes the method's application in detail to an upcoming formation flying mission.

A. Reference Case

A collision probability test case from Alfano[15] is adapted to demonstrate the WSPRT. This case, number three in Alfano's paper, involves a close approach of 3.9 m between two objects in geosynchronous Earth orbit (GEO), whose combined hard-body radius is 15 m. The conjunction occurs with high enough relative velocity, 16 m/s, that collision probability can be accurately computed using an integral over the uncertainty projected into a plane normal to the relative velocity at the time of closest approach. Alfano gives the states and covariances of the two objects at the time of closest approach, and at an epoch 3.25 (solar) days prior, and states that he used two-body dynamics to propagate the objects. Alfano computes the collision probability using a variety of methods, and finds that two-dimensional linear approaches produce $P_c = 10.035\%$, which differs from the relative frequency of collision computed from 1×10^8 Monte Carlo trials by only 0.49%.

For purposes of this demonstration, relevant portions of Figure 4 were simulated. In particular, Alfano's Case 3 relative state and covariance were interpreted as the prior statistics required by Steps 1a and 1b, and hence the background rate of collision computed in Step 1c was the collision probability computed by Alfano; that is, $P_{c|o} = 10.035\%$. These statistics were used to generate Monte Carlo samples of the true miss distance at the time of closest approach, so that about 10% of the trials resulted in hits. Each sample seeded a trial of the WSPRT, and the experiment would be expected to demonstrate that the WSPRT properly detects most of the resulting hits, with missed detection and false alarm rates consistent with the target rates chosen in Step 1d, which were taken as $\bar{P}_{fa} = 20\%$ and $\bar{P}_{md} = 1\%$. Consistent with prior studies,[1,2] the orbit determination and tracking processes of Steps 2 and 3 were not simulated. Instead, corresponding to each true miss distance, a sequence of noisy predicted miss vectors and associated formal covariances were generated at intervals of

Table 3. Predictive position covariance data at time of closest approach expressed in radial, in-track, cross-track axes defined by object 2's orbital state. The terms above the diagonals are covariances in square meters, the terms below the diagonals are correlation coefficients, and the terms along the diagonal are standard deviations in meters. A significant decrease of in-track position uncertainty as the prediction time decreases is evident.

$\tilde{\boldsymbol{P}}_{*|1}$	$(t_1 = t_* - 3.25 \text{ days})$	
$\sigma_r = 6.11221$	$p_{ri} = -682.918$	$p_{rc} = 1.30493$
$\rho_{ir} = -0.978472$	$\sigma_i = 114.188$	$p_{ic} = -24.0516$
$\rho_{cr} = 0.135791$	$\rho_{ci} = -0.13397$	$\sigma_c = 1.57223$
$\tilde{\boldsymbol{P}}_{*|2}$	$(t_2 = t_* - 2.25 \text{ days})$	
$\sigma_r = 5.61564$	$p_{ri} = -416.266$	$p_{rc} = 0.787862$
$\rho_{ir} = -0.969833$	$\sigma_i = 76.4319$	$p_{ic} = -10.7667$
$\rho_{cr} = 0.0805311$	$\rho_{ci} = -0.0808579$	$\sigma_c = 1.74216$
$\tilde{\boldsymbol{P}}_{*|3}$	$(t_3 = t_* - 1.25 \text{ days})$	
$\sigma_r = 5.06425$	$p_{ri} = -191.592$	$p_{rc} = 0.358523$
$\rho_{ir} = -0.967015$	$\sigma_i = 39.1228$	$p_{ic} = -2.81429$
$\rho_{cr} = 0.0381662$	$\rho_{ci} = -0.0387807$	$\sigma_c = 1.85491$
$\tilde{\boldsymbol{P}}_{*|4}$	$(t_4 = t_* - 0.25 \text{ days})$	
$\sigma_r = 4.47697$	$p_{ri} = -14.6702$	$p_{rc} = 0.0270485$
$\rho_{ir} = -0.748294$	$\sigma_i = 4.37904$	$p_{ic} = -0.0285942$
$\rho_{cr} = 0.00317455$	$\rho_{ci} = -0.00343102$	$\sigma_c = 1.90317$

one day, starting at 3.25 days prior to closest approach and ending at 0.25 days prior. These predictions correspond to Step 4 in Figure 4. To generate each set of predictions, it was assumed that an independent epoch state solution for the orbits of the two objects would be available at each of the times of interest, and that the covariance of this epoch solution was the same as Alfano's epoch state covariance. Then two-body dynamics propagated Alfano's epoch covariances to the close approach time. The noisy predictions were then produced by drawing samples from a Gaussian whose mean was the true miss vector for that trial. Table 3 expresses the position portions of the predicted covariances computed in this manner, in a radial, in-track, cross-track coordinate frame defined by the orbit of object 2.

It follows from the description of the simulation procedure that the probabilities defined in Eq. (13) can be computed using the procedure that the appendix describes. Table 4 shows the position portions of the total covariances computed accordingly. A comparison of Table 4 with Table 3 illustrates the improvement in accuracy that results from accumulating information. Given the target false alarm rate of $\bar{P}_{fa} = 20\%$ and target missed detection rate of $\bar{P}_{md} = 1\%$, and that $P_{c|o} \approx 10\%$, the Monte Carlo method would be expected to need about 10,000 trials in order to converge to repeatable results. Out of the 10,000 trials that were
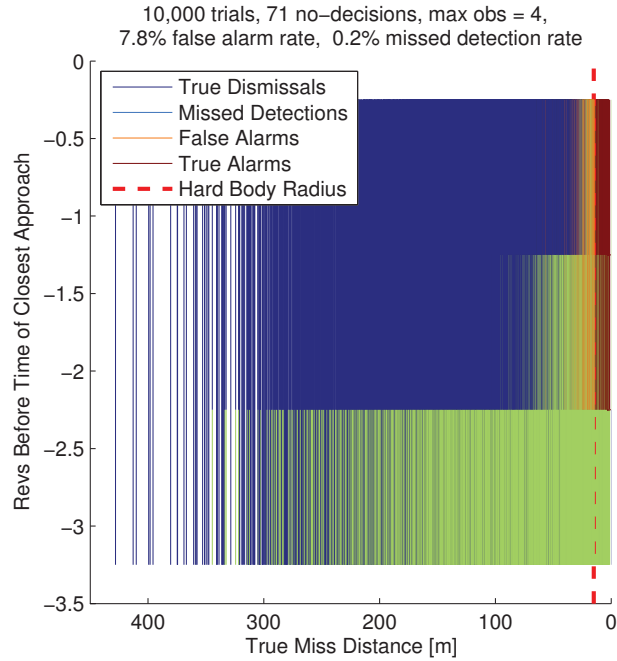
Figure 5. GEO Example (Alfano Test Case 3). Each vertical line corresponds to a time series of $P_{c|k}$ for one Monte Carlo trial, with earliest time at the bottom and time of closest approach at top. A dark blue color indicates $P_{c|k} < P_c^D$ when a miss occurs; light blue indicates $P_{c|k} < P_c^D$ when a hit occurs; orange indicates $P_{c|k} \geq P_c^A$ when a miss occurs; dark red indicates $P_{c|k} \geq P_c^A$ when a hit occurs; and green indicates $P_c^D \leq P_{c|k} < P_c^A$. Two missed detections that occurred at true miss distances of 8 and 10.5 meters are obscured by the closely-packed plot lines; in both of these cases $P_{c|1} < P_c^D$.

Table 4. Total position covariance, computed according to Eq. (23), at time of closest approach, expressed as in Table 3.

$\hat{\boldsymbol{P}}_{*|1}$		$(t_1 = t_* - 3.25 \text{ days})$	
	$\sigma_r = 4.32198$	$p_{ri} = -341.459$	$p_{rc} = 0.652464$
	$\rho_{ir} = -0.978472$	$\sigma_i = 80.7434$	$p_{ic} = -12.0258$
	$\rho_{cr} = 0.135791$	$\rho_{ci} = -0.13397$	$\sigma_c = 1.11174$
$\hat{\boldsymbol{P}}_{*|2}$		$(t_2 = t_* - 2.25 \text{ days})$	
	$\sigma_r = 1.01314$	$p_{ri} = -11.4757$	$p_{rc} = 0.0912979$
	$\rho_{ir} = -0.740403$	$\sigma_i = 15.2982$	$p_{ic} = -1.88144$
	$\rho_{cr} = 0.110861$	$\rho_{ci} = -0.1513$	$\sigma_c = 0.812854$
$\hat{\boldsymbol{P}}_{*|3}$		$(t_3 = t_* - 1.25 \text{ days})$	
	$\sigma_r = 0.590341$	$p_{ri} = -1.1996$	$p_{rc} = 0.0103076$
	$\rho_{ir} = -0.548492$	$\sigma_i = 3.70478$	$p_{ic} = -0.352589$
	$\rho_{cr} = 0.0294343$	$\rho_{ci} = -0.160438$	$\sigma_c = 0.593198$
$\hat{\boldsymbol{P}}_{*|4}$		$(t_4 = t_* - 0.25 \text{ days})$	
	$\sigma_r = 0.457077$	$p_{ri} = -0.571993$	$p_{rc} = -0.000977275$
	$\rho_{ir} = -0.949736$	$\sigma_i = 1.31765$	$p_{ic} = 0.00786831$
	$\rho_{cr} = -0.00480251$	$\rho_{ci} = 0.0134129$	$\sigma_c = 0.445204$

simulated, 978 resulted in a close approach of less than the 15 m combined hard-body radius, of which two failed to be detected, for an achieved missed detection rate of $P_{md} = 0.20\%$. In both of these trials, the dismissal occurred on the first observation, which suggests that in practice, the dismissal would have ample opportunity to be reversed when additional data arrived. Of the 9,022 misses, alarms were reported in 705 trials, for an achieved false alarm rate of $P_{fa} = 7.81\%$. In 71 trials, no decision was reached at the data cutoff time of -0.25 days. The miss distances for these no-decision cases ranged from 13.6 m to 21.6 m. Out of these 71, seven cases corresponded to hits. Assuming in practice that prudence would dictate an alarm in such no-decision cases, the 64 that were misses would increase the false alarm rate to $P_{fa} = 8.52\%$. Figure 5 illustrates these results. The figure shows that large misses are dismissed immediately at -3.25 days with great frequency. Almost all of the misses greater than 100 m are dismissed after an additional observation at -2.25 days, as are the majority of decisions of any type. All of the false alarms occur for miss distances between 65 m and the 15 m hard-body radius.

B. Operational Case

A more realistic test case is adapted from Foster's description of ISS debris avoidance operations.[1] In this work, Foster modeled the background debris flux in the ISS orbit circa 1997, finding it to be approximately 0.27 objects per square kilometer of projected area per year. With a 60 square meter average area projected into the conjunction plane, this flux produces a background probability of collision for the ISS of 0.0031. Foster shows how selecting a single maneuver/dismissal threshold on collision probability of 10^{-4} reduces the overall risk by 80%, while requiring an average of two risk mitigation maneuvers per year.

To demonstrate the enhancement the WSPRT might offer for this case, take $P_{c|o} = 0.0031$, and set the dismissal threshold $P_c^D = 10^{-4}$, so as to make minimal change to existing procedure. Also assume that if $P_{c|k}$ exceeds a threshold of $P_c^A = 10^{-2}$, a maneuver would be warranted under almost any circumstances, given the prominence of the ISS. Eqs. (19) and (20) may be inverted to show that these thresholds correspond to target false alarm and missed detection rates of $\bar{P}_{fa} = 29\%$ and $\bar{P}_{md} = 2.4\%$, respectively. The former implies that there would be a 29% probability that if the ISS performs a debris avoidance maneuver, it was not actually necessary to avoid being hit be a debris object. The later implies that the fractional residual risk, i.e. the risk that a debris strike will occur when $P_{c|k} \leq 10^{-4}$, will be reduced from Foster's value of 20% to less than 3%.

Although Foster performed a study of orbit prediction accuracy in Reference 1, the results are scaled in such a way that makes their numerical values undecipherable. However, based on Figure 6 in that work, one can approximately determine standard deviations of a Gaussian error distribution in the conjunction plane that will reproduce $P_{c|o} = 0.0031$ to sufficient accuracy for present purposes. Foster's Figure 6 places the 60 m hard body disk at its origin, and depicts elliptical contours of constant relative debris position probability. Taking the vertical standard deviation to be 450 m, and the horizontal standard deviation to be 1300 m, and integrating over the 60 m hard body disk at the origin, results in $P_{c|o} = 0.0030693$. To demonstrate the WSPRT, these standard deviations were taken as the as the prior statistics required by Steps 1a and 1b of Figure 4, and were used to generate Monte Carlo samples of the true miss distance at the time of closest approach, so that about 0.3% of the trials resulted in hits. As with the previous example, each sample seeded a trial of the WSPRT, and the experiment would be expected to demonstrate that the WSPRT properly detects most of the resulting hits, with missed detection and false alarm rates consistent with the target rates of $\bar{P}_{fa} = 29\%$ and $\bar{P}_{md} = 2.4\%$, determined as described above.

Since Foster's prediction statistics are not available, the uncertainty prediction model developed by Jenkin[2] is adopted. Jenkin states that his model is representative of the prediction accuracy available to the robotic spaceflight community, which is less accurate than the data provided for use by the ISS. As with the previous example, the orbit determination

and tracking processes of Steps 2 and 3 were not simulated. Instead, corresponding to each true miss distance, a sequence of noisy predicted miss vectors and associated formal covariances were generated at intervals of one day, starting at 3 days prior to closest approach and ending at 8 hours prior, corresponding to Step 4 in Figure 4. As before, it was assumed that an independent epoch state solution for the orbits of the two objects would be available at each of the times of interest. The covariance corresponding to these solutions, predicted to the time of closest approach, was taken from Jenkin's model. In particular, Jenkin's radial standard deviations are taken as vertical conjunction plane standard deviations, and his normal standard deviations are taken as horizontal conjunction plane standard deviations (where "vertical" and "horizontal" correspond to Foster's Figure 6 notation); Jenkin's standard deviations for "cooperatively ranged [low earth orbit] satellites" are taken for the primary object (ISS), and his standard deviations for "typical secondary objects" are taken for the debris. The corresponding variances were then summed to give the (diagonal) relative state covariance predictions, which were then used to draw the noisy predictions for each trial, where the mean was the true miss vector for that trial. Figure 6 shows the time series of resulting predictive relative state standard deviations at the time of closest approach, plotted with the *a priori* standard deviations discussed in the previous paragraph. The figure also shows the total position covariance, computed according to Eq. (23). The plots therefore correspond to Tables 3 and 4 from the previous example.
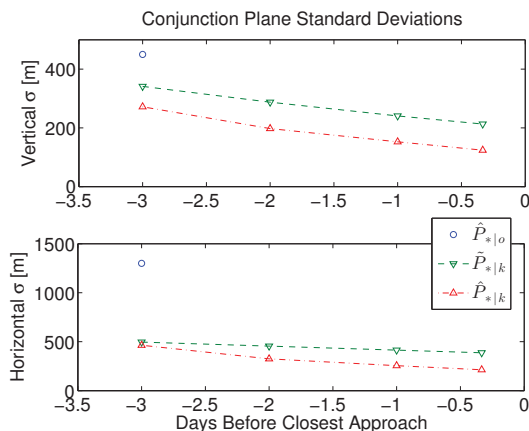
Figure 6. Predictive uncertainties at time of closest approach for the ISS Example, adapted from Jenkin.[2]

Given the target false alarm rate of $\bar{P}_{fa} = 29\%$ and target missed detection rate of $\bar{P}_{md} = 2.4\%$, and that $P_{c|o} = 0.31\%$, the Monte Carlo method would be expected to need about 1,000,000 trials in order to converge to repeatable results. Out of the 1,000,000 trials that were simulated, 3084 resulted in a close approach of less than the 60 m hard-body radius, of which 23 failed to be detected, for an achieved missed detection rate of $P_{md} = 0.75\%$. Of the 996,916 misses, alarms were reported in 259,910 trials, for an achieved false alarm rate of

$P_{fa} = 26.07\%$. In 62,337 trials, no decision was reached at the data cutoff time of -8 hours. Out these 62,337, 91 cases actually resulted in hits. If we assume that a maneuver would be performed in all such no-decision cases, then there would be 62,245 additional false alarms, increasing the effective false alarm rate to 32.3%. Figure 7 illustrates these results. As the figure shows, the ISS dismissal threshold of 10^{-4} results in a large range of miss distances, up to 4383 m for this demonstration, that correspond to false alarms.
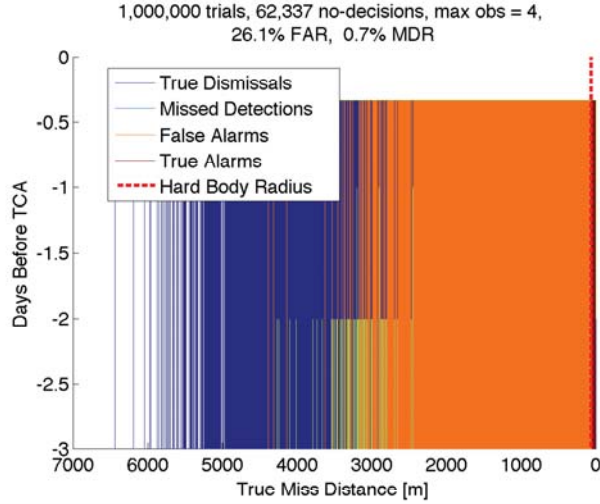

Figure 7. ISS Example. Each vertical line corresponds to a time series of $P_{c|k}$ for one Monte Carlo trial, with earliest time at the bottom and time of closest approach at top. A dark blue color indicates $P_{c|k} < P_c^D$ when a miss occurs; light blue indicates $P_{c|k} < P_c^D$ when a hit occurs; orange indicates $P_{c|k} \geq P_c^A$ when a miss occurs; dark red indicates $P_{c|k} \geq P_c^A$ when a hit occurs; and green indicates $P_c^D \leq P_{c|k} < P_c^A$.

VII. Conclusions

The Wald Sequential Probability Ratio Test (WSPRT) developed in this work should provide a remarkably simple method for owner/operators of maneuverable spacecraft to quantify their conjunction risk mitigation maneuver decisions. This quantification involves a tolerance for false alarms, a tolerance for missed detections, and an assessment of the basic risk of collision for a given encounter. Although in general a WSPRT does not terminate within a definite time, the examples analyzed in the paper showed that for test cases available from the relevant literature, the WSPRT method yields expected results within time horizons that appear to be suitable for an operational conjunction assessment decision timeline.

Notably, this work has highlighted the significant role that prior perception of risk, quantified herein as the probability $P_{c|o}$, plays in the decision process. It has shown that missed detection and false alarm rates must be understood in the context of the underlying risk that $P_{c|o}$ represents, in order to properly understand the outcome of a decision procedure. Prior

work on conjunction assessment has primarily treated uncertainty as epistemic, arising from noise in an observation process, which the present work quantifies as $P_{c|k}$. The dependence of the present result on *both* $P_{c|o}$ and $P_{c|k}$ clearly shows that aleatoric uncertainty clouds the decision process, independent of any error of observation.

An important contribution of this work is to show that a compound hypothesis test concerning the conditioning of a set of observations, $p_{\mathbf{Y}_k}(\mathbb{Y}_k \,|\, \mathbf{x}_t \in \mathbb{U}_I)$ in the case of collision avoidance, which is normally unavailable, reduces to a ratio test involving $p_{\mathbf{x}_t}(\boldsymbol{\xi}_t \,|\, \mathbb{Y}_k)$. In the Gaussian case, inclusive of Gaussian mixtures, a Bayesian computation of $p_{\mathbf{x}_t}(\boldsymbol{\xi}_t \,|\, \mathbb{Y}_k)$ is readily available. This result does not appear in Wald's book,[4] and the authors have not found it elsewhere in the literature.

A consideration of Bayesian estimators such as the Kalman filter is that they generally allow for the possibility of additive process noise. Designers often introduce such noise as an artificial tuning measure. However, it is important to emphasize that if such process noise is actually present, and gives rise to disturbances to the state whose integrated effect could allow the system to change from the null to the alternative hypothesis (that is, change from a hit to a miss), then a WSPRT is not appropriate. Since presently available methods of computing collision probability also exclude the possibility of such changes, in that they assume all uncertainty arises from initial conditions, this does not appear to be a limitation of the decision procedure proposed in this paper.

The simplicity of the WSPRT procedure belies the complexity of computing and then integrating $p_{\mathbf{x}_t}(\boldsymbol{\xi}_t \,|\, \mathbb{Y}_k)$ for the conjunction assessment scenario, which remains a challenging problem, despite numerous recent advances. To the extent that inadequately characterized error statistics, biased state estimation and inadequate computation of collision probability hamper the current conjunction assessment process, the method proposed here might be expected to suffer as well.

Appendix: Gaussian Case

Assume the true relative position has a Gaussian distribution, with some fixed but unknown mean and covariance:

$$p_{\boldsymbol{r}_*}(\boldsymbol{r}_*) = \mathrm{N}\left(\boldsymbol{r}_* | \boldsymbol{\mu}_*, \boldsymbol{P}_*\right) = \frac{1}{(2\pi)^{\frac{n}{2}}\sqrt{|\boldsymbol{P}_*|}} \exp\left(-\frac{1}{2}(\boldsymbol{r}_* - \boldsymbol{\mu}_*)^{\top}\boldsymbol{P}_*^{-1}(\boldsymbol{r}_* - \boldsymbol{\mu}_*)\right) \qquad (21)$$

where n is the dimension of the relative position. Assume the predicted relative positions have Gaussian distributions, centered on the true relative position, so the joint density, after

k observations, becomes

$$
\mathrm{p}_{\tilde{\boldsymbol{r}}_{*|k}, \tilde{\boldsymbol{r}}_{*|k-1}, \ldots, \tilde{\boldsymbol{r}}_{*|1}, \boldsymbol{r}_*} \left(\tilde{\boldsymbol{r}}_{*|k}, \tilde{\boldsymbol{r}}_{*|k-1}, \ldots, \tilde{\boldsymbol{r}}_{*|1}, \boldsymbol{r}_* \right) = \prod_{i=1}^{k} \left(\frac{1}{(2\pi)^{\frac{n}{2}} \sqrt{|\tilde{\boldsymbol{P}}_{*|i}|}} \right) \frac{1}{(2\pi)^{\frac{n}{2}} \sqrt{|\boldsymbol{P}_*|}}
$$
$$
\cdot \exp\left(-\frac{1}{2} \left[\sum_{i=1}^{k} (\tilde{\boldsymbol{r}}_{*|i} - \boldsymbol{r}_*)^\top \tilde{\boldsymbol{P}}_{*|i}^{-1} (\tilde{\boldsymbol{r}}_{*|i} - \boldsymbol{r}_*) + (\boldsymbol{r}_* - \boldsymbol{\mu}_*)^\top \boldsymbol{P}_*^{-1} (\boldsymbol{r}_* - \boldsymbol{\mu}_*) \right] \right)
\tag{22}
$$

Since $\boldsymbol{\mu}_*$ and $\boldsymbol{P}_*$ are unknown, *a priori* estimates $\hat{\boldsymbol{r}}_{*|o}$ and $\hat{\boldsymbol{P}}_{*|o}$ are used instead. Collecting terms as follows will result in a considerable simplification.

$$
\hat{\boldsymbol{P}}_{*|k} = \left(\hat{\boldsymbol{P}}_{*|o}^{-1} + \sum_{i=1}^{k} \tilde{\boldsymbol{P}}_{*|i}^{-1} \right)^{-1}
\tag{23}
$$

$$
\hat{\boldsymbol{r}}_{*|k} = \hat{\boldsymbol{P}}_{*|k} \left(\hat{\boldsymbol{P}}_{*|o}^{-1} \hat{\boldsymbol{r}}_{*|o} + \sum_{i=1}^{k} \tilde{\boldsymbol{P}}_{*|i}^{-1} \tilde{\boldsymbol{r}}_{*|i} \right)
\tag{24}
$$

$$
\beta = \hat{\boldsymbol{r}}_{*|o}^\top \hat{\boldsymbol{P}}_{*|o}^{-1} \hat{\boldsymbol{r}}_{*|o} + \sum_{i=1}^{k} \tilde{\boldsymbol{r}}_{*|i}^\top \tilde{\boldsymbol{P}}_{*|i}^{-1} \tilde{\boldsymbol{r}}_{*|i} - \hat{\boldsymbol{r}}_{*|k}^\top \hat{\boldsymbol{P}}_{*|k}^{-1} \hat{\boldsymbol{r}}_{*|k}
\tag{25}
$$

So the conditional densities associated with the two hypotheses become

$$
\mathrm{p}_{\mathbf{Y}_k}(\mathbb{Y}_k \,|\, \boldsymbol{r}_* \in \mathbb{U}_I) = \frac{\prod_{i=1}^{k} \left(\frac{1}{(2\pi)^{\frac{n}{2}} \sqrt{|\tilde{\boldsymbol{P}}_{*|i}|}} \right) \sqrt{\frac{|\hat{\boldsymbol{P}}_{*|k}|}{|\hat{\boldsymbol{P}}_{*|o}|}} \, \mathrm{e}^{-\frac{1}{2}\beta} \, P_{c|k}}{P_{c|o}}
\tag{26}
$$

$$
\mathrm{p}_{\mathbf{Y}_k}(\mathbb{Y}_k \,|\, \boldsymbol{r}_* \notin \mathbb{U}_I) = \frac{\prod_{i=1}^{k} \left(\frac{1}{(2\pi)^{\frac{n}{2}} \sqrt{|\tilde{\boldsymbol{P}}_{*|i}|}} \right) \sqrt{\frac{|\hat{\boldsymbol{P}}_{*|k}|}{|\hat{\boldsymbol{P}}_{*|o}|}} \, \mathrm{e}^{-\frac{1}{2}\beta}(1 - P_{c|k})}{1 - P_{c|o}}
\tag{27}
$$

where

$$
P_{c|k} = \frac{1}{(2\pi)^{\frac{n}{2}} \sqrt{|\hat{\boldsymbol{P}}_{*|k}|}} \int_{\mathbb{U}_I} \exp\left(-\frac{1}{2} (\hat{\boldsymbol{r}}_{*|k} - \boldsymbol{r}_*)^\top \hat{\boldsymbol{P}}_{*|k}^{-1} (\hat{\boldsymbol{r}}_{*|k} - \boldsymbol{r}_*) \right) \mathrm{d}\boldsymbol{r}_*
\tag{28}
$$

which can be defined as the instantaneous collision probability, and

$$
P_{c|o} = \frac{1}{(2\pi)^{\frac{n}{2}} \sqrt{|\hat{\boldsymbol{P}}_{*|o}|}} \int_{\mathbb{U}_I} \exp\left(-\frac{1}{2} (\boldsymbol{r}_* - \hat{\boldsymbol{r}}_{*|o})^\top \hat{\boldsymbol{P}}_{*|o}^{-1} (\boldsymbol{r}_* - \hat{\boldsymbol{r}}_{*|o}) \right) \mathrm{d}\boldsymbol{r}_*
\tag{29}
$$

which can be viewed as a prior estimate of instantaneous collision probability, determined from considerations independent of the data that produced the predictions presently avail-

able. Dividing Eq. (27) by Eq. (26) gives Eq. (14).

Acknowledgments

The authors wish to thank Donald J. Dichmann and Trevor W. Williams of the Navigation and Mission Design Branch at Goddard Space Flight Center for reviewing this work.

References

[1] Foster, Jr., J. L., "The Analytical Basis for Debris Avoidance Operations for the International Space Station," *Third European Conference on Space Debris*, ESA Publications Division of the European Space Research and Technology Centre, Noordwijk, The Netherlands, 1997, pp. 441–445.

[2] Jenkin, A. B., "Effect of Orbit Data Quality on the Feasibility of Collision Risk Management," *Journal of Spacecraft and Rockets*, Vol. 41, No. 4, 2014/02/24 2004, pp. 677–683, doi: 10.2514/1.11941.

[3] Patera, R. P. and Peterson, G. E., "Space Vehicle Maneuver Method to Lower Collision Risk to an Acceptable Level," *Journal of Guidance, Control, and Dynamics*, Vol. 26, No. 2, 2014/02/24 2003, pp. 233–237, doi: 10.2514/2.5063.

[4] Wald, A., *Sequential Analysis*, chap. 1–3, Dover Publications, 2004, isbn:0486439129.

[5] Foster, Jr., J. L. and Estes, H. S., "A Parametric Analysis of Orbital Debris Collision Probability and Maneuver Rate for Space Vehicles," Tech. Rep. JSC–25898, NASA Johnson Space Center, Houston, TX, 1992.

[6] Akella, M. R. and Alfriend, K. T., "Probability of Collision Between Space Objects," *Journal of Guidance, Control, and Dynamics*, Vol. 23, No. 5, 2000, pp. 769–772, doi: 10.2514/2.4611.

[7] Patera, R. P., "General Method for Calculating Satellite Collision Probability," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 4, 2001, pp. 716–722, doi: 10.2514/2.4771.

[8] Chan, K., "Improved Analytical Expressions for Computing Spacecraft Collision Probabilities," *Space Flight Mechanics 2003*, Vol. 114 of *Advances in the Astronautical Sciences*, Univelt, 2003.

[9] Alfano, S., "A Numerical Implementation of Spherical Object Collision Probability," *Journal of the Astronautical Sciences*, Vol. 53, No. 1, Jan–Mar 2005, pp. 103–109.

[10] Alfano, S., "Review of Conjunction Probability Methods for Short-Term Encounters," *Astrodynamics 2007*, Vol. 129 of *Advances in the Astronautical Sciences*, Univelt, 2008.

[11] Patera, R. P., "Satellite Collision Probability for Nonlinear Relative Motion," *Journal of Guidance, Control, and Dynamics*, Vol. 26, No. 5, 2003, pp. 728–733, doi: 10.2514/2.5127.

[12] Chan, K., "Spacecraft Collision Probability for Long-Term Encounters," *Astrodynamics 2003*, Vol. 116 of *Advances in the Astronautical Sciences*, Univelt, 2004.

[13] McKinley, D., *Development of a Nonlinear Probability of Collision Tool for the Earth Observing System*, American Institute of Aeronautics and Astronautics, 2006, doi:10.2514/6.2006-6295.

[14] Alfano, S., *Addressing Nonlinear Relative Motion For Spacecraft Collision Probability*, American Institute of Aeronautics and Astronautics, 2006, doi:10.2514/6.2006-6760.

[15] Alfano, S., "Satellite Conjunction Monte Carlo Analysis," *Space Flight Mechanics 2009*, Vol. 135 of *Advances in the Astronautical Sciences*, Univelt, San Diego, CA, 2010.

[16]DeMars, K. J., Bishop, R. H., and Jah, M. K., "A Splitting Gaussian Mixture Method for the Propagation of Uncertainty in Orbital Mechanics," *Space Flight Mechanics 2011*, Vol. 140 of *Advances in the Astronautical Sciences*, Univelt, 2012.

[17]Jones, B. A., Doostan, A., and Born, G. H., "Nonlinear Propagation of Orbit Uncertainty Using Non-Intrusive Polynomial Chaos," *Journal of Guidance, Control and Dynamics*, Vol. 36, No. 2, March 2013, pp. 430–444.

[18]Coppola, V. T., "Including Velocity Uncertainty in the Probability of Collision Between Space Objects," *Space Flight Mechanics 2012*, Vol. 143 of *Advances in the Astronautical Sciences*, Univelt, 2012.

[19]Coppola, V. T., "Evaluating the Short-Term Encounter Assumption of the Probability of Collision Formula," *Space Flight Mechanics 2012*, Vol. 143 of *Advances in the Astronautical Sciences*, Univelt, 2012.

[20]Carpenter, J. R., "Non-Parametric Collision Probability for Low-Velocity Encounters," *Space Flight Mechanics 2007*, Vol. 127 of *Advances in the Astronautical Sciences*, Univelt, 2007.

[21]Kumar, M., Singla, P., Chakravorty, S., and Junkins, J., *The Partition of Unity Finite Element Approach to the Stationary Fokker-Planck Equation*, American Institute of Aeronautics and Astronautics, 2006, doi:10.2514/6.2006-6285.

[22]Shiryaev, A. N., *Optimal Stopping Rules*, Vol. 8 of *Applications of Mathematics*, chap. 4, Springer, 1st ed., 1978, isbn:3540740104.

[23]DeMars, K. J., Cheng, Y., and Jah, M. K., "Collision Probability with Gaussian Mixture Orbit Uncertainty," *Journal of Guidance, Control, and Dynamics*, 2014, pp. 1–6, doi: 10.2514/1.62308.

[24]Wawrzyniak, G. G., Carpenter, J. R., Mattern, D. J., Williams, T. W., and Ottenstein, N. A., "Conjunction Assessment Concept Of Operations For The Magnetospheric Multi-Scale (MMS) Mission," *Astrodynamics 2013*, Vol. 150 of *Advances in the Astronautical Sciences*, Univelt, 2013.

[25]Carpenter, J. R., Markley, F. L., and Gold, D., "Sequential Probability Ratio Test for Collision Avoidance Maneuver Decisions," *Journal of the Astronautical Sciences*, Vol. 59, No. 1,2, January–June 2012, pp. 273–286.

[26]Carpenter, J. R., Markley, F. L., Alfriend, K. T., Wright, C., and Arcido, J., "Sequential Probability Ratio Test for Collision Avoidance Maneuver Decisions Based on a Bank of Norm-Inequality-Constrained Epoch-State Filters," *Astrodynamics 2011*, Vol. 142 of *Advances in the Astronautical Sciences*, Univelt, 2012.

[27]Carpenter, J. R. and Markley, F. L., "Current-State Constrained Filter Bank for Wald Testing of Spacecraft Conjunctions," *Proceedings of the 23rd International Symposium on Space Flight Dynamics*, www.issfd.org, 2012.

[28]Carpenter, J. R., Markley, F. L., and Gold, D., *Wald Sequential Probability Ratio Test for Analysis of Orbital Conjunction Data*, American Institute of Aeronautics and Astronautics, 2013, doi:10.2514/6.2013-5187.

[29]Patera, R. P., "Space Vehicle Conflict Probability for Ellipsoidal Conflict Volumes," *Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 6, 2007, pp. 1819–1822, doi: 10.2514/1.30504.

[30]Maybeck, P. S., *Stochastic Models, Estimation, and Control*, Vol. 1, chap. 3, Academic Press, New York, 1979, isbn:0124807011.

[31]Newman, L. K., "The NASA Robotic Conjunction Assessement Process: Overview and Operational Experience," *59th International Astronautical Congress*, Vol. 6, International Astronautical Federation, Glasgow, Scotland, 29 September – 03 October 2008, pp. 2260–2274.

[32]McKinley, D., Rohrbaugh, D., and Carpenter, R., "Conjunction Assessment for the Magnetospheric Multi-Scale Formation," *Astrodynamics 2009*, Vol. 135 of *Advances in the Astronautical Sciences*, Univelt, 2010.

[33]Narvet, S., Frigm, R. C., and Hejduk, M., "Assessment of Uncertainty-Based Screening Volumes for NASA Robotic LEO and GEO Conjunction Risk Assessment," *Astrodynamics 2011*, Vol. 142 of *Advances in the Astronautical Sciences*, Univelt, 2011.