

Sequential Wald Test Employing a Norm-Inequality-Constrained Filter Bank

with Application to Spacecraft Conjunctions

J. Russell Carpenter · F. Landis Markley

Received: 23 Sep 2020 / Accepted: 12 Mar 2021
Communicated by Jason L. Speyer

Abstract A binary Wald sequential probability ratio test that uses the residuals of two norm-inequality-constrained Kalman filters for its likelihood ratio is employed for a class of compound hypothesis tests on non-stationary systems. The hypotheses concern an inequality constraint on the norm of some elements of the system state. Each of the two constrained Kalman filters minimizes the summed-squares of its estimation errors subject to one or the other direction of the inequality constraint. This test is applied to the problem of spacecraft conjunction assessment, wherein the constraint concerns the close approach distance between a spacecraft and another space object. The outcome of the test can inform decisions concerning risk mitigation maneuvers by an active spacecraft.

Keywords constrained Kalman filtering · decision theory · spacecraft conjunction analysis

Mathematics Subject Classification (2000) 62M99 · 62C99 · 70F99

1 Introduction

Wald's sequential probability ratio test [20] (WSPRT) forms the basis for much of modern decision theory. In contrast to prior work, Wald showed that constraints on both false alarm and missed detection rates can be asymptotically

Russell Carpenter, Corresponding author
Goddard Space Flight Center
Greenbelt, Maryland
Russell.Carpenter@nasa.gov

Landis Markley
Retired
South Hadley, Massachusetts
Landis.Markley@comcast.net

met by sequentially updating a likelihood ratio formed from successive sample evaluations of probability densities conditioned on each of a pair of mutually exclusive, fixed hypotheses. In its most common form, the Wald test is applied to tests of simple hypothesis concerning stationary systems. A classic example is testing a hypothesis on the mean of a set of samples purported to be drawn from a normal distribution. Many real-world problems instead are best modeled by non-stationary systems, and many real-world decisions concern continua of hypotheses, known as a compound hypotheses [11]. Note that Wald's SPRT depends on the assumption that the hypotheses being compared are fixed. If one wishes to consider that random or unmodeled perturbations could change characteristics of the hypotheses during the course of testing, then a different test, such as a Shiryaev test [19], is appropriate.

The contribution of the present work is to show how the residuals of a bank [12] of norm-inequality-constrained Kalman filters can form the likelihood ratio for a WSPRT that concerns compound hypotheses on a non-stationary system, motivated by the spacecraft collision avoidance decision problem [1]. The competing compound hypotheses that the WSPRT adjudicates take the form of a complementary pair of inequality constraints on the norm of some elements of the system state at a specified future time. The inequality-constrained filters are novel extensions to work done on equality-constrained filtering [10,21]. The specified time, t_* is strictly posterior to the collection of the measurement data processed by the filters. Because the constraint applies only at the future time t_* , the filters may either take the form of epoch-state estimators [14,13], where the epoch is chosen to be t_* , or take the usual “current-state” form of the Kalman filter, if the future-time constraint is mapped from t_* to each filter measurement update time. To enforce common measurement editing histories, and optionally to provide a linearization reference, if desired or needed, a third unconstrained Kalman filter is loosely coupled to the processing architecture. Fig. 1 provides an overview of the filter bank. The authors have also developed a different approach to a version of this




Fig. 1 Filter Bank for WSPRT. The set of measurements $\mathbb{Y}_{1:k}$ is processed by two inequality-constrained Kalman filters, and one unconstrained Kalman filter. The conditional densities of the innovations of the constrained filters, $\bar{\epsilon}_{k|\mathcal{H}_i}$, are used in a likelihood ratio, Λ_k , for the WSPRT.

problem, which does not involve filter banks [6]. In that work, the conditional densities required for the WSPRT likelihood ratio are reduced to integrals of the densities derived from a single, unconstrained estimator, over the domains of the two competing hypotheses. Accurate and tractable calculation of such integrals may not always be feasible however, so the approach of the present work remains of interest.

The remainder of this work proceeds as follows. The next section concisely states the general problem, and the subsequent section provides two forms of a solution: the first uses epoch-state Kalman filters, while the second uses more common current-state filters. Next, a discussion compares and contrasts the two solution approaches. A penultimate section illustrates the application of both solution approaches to the topic of spacecraft conjunctions. The final section offers conclusions.

2 Problem Statement

Let the time-evolution of the state of a system of interest occur according to the stochastic differential equation

$$d\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t), t) dt + \mathbf{G}(t) d\mathbf{w}(t), \quad (1)$$

where $\mathbf{x} \in \mathbb{R}_n$, $\mathbf{w} \in \mathbb{R}_p$, $p \leq n$, with initial condition $\mathbf{x}(t_o) = \mathbf{x}_o$. The process noise input $d\mathbf{w}(t)$ is a zero-mean Gaussian, time-wise uncorrelated stochastic process with power spectral density $\mathbf{Q}$. An estimate of the state at the initial time is $\hat{\mathbf{x}}_o$. The error of this estimate, $\mathbf{e}_o = \mathbf{x}_o - \hat{\mathbf{x}}_o$, follows a Gaussian distribution with zero mean and covariance $\mathbf{P}_o$. Measurements of the system occur as a discrete-time sequence of finite length according to

$$\mathbf{y}(t_k) = \mathbf{h}(\mathbf{x}(t_k), t_k) + \mathbf{v}(t_k), \quad (2)$$

where $\mathbf{y} \in \mathbb{R}_m$, $m \leq n$, $k \in \{1, 2, \dots, k_f\}$, and the measurement noise $\mathbf{v} \in \mathbb{R}_m$ is a zero-mean, Gaussian, time-wise uncorrelated stochastic sequence with covariance $\mathbf{R}(t_k)$. The initial state error is uncorrelated for all time with the measurement noise and process noise, and the measurement noise and process noise are also uncorrelated with each other.

At any time inclusively prior to t_{k_f} , the problem is to choose between two complementary compound hypotheses concerning a two-normed¹ subset of the state, $\mathbf{s} \in \mathbb{R}_q$, $q \leq n$, at the specified time $t_* > t_{k_f}$:

$$\mathcal{H}_0 : \|\mathbf{s}(t_*)\| \leq \mathcal{S} \quad (3)$$

$$\mathcal{H}_1 : \|\mathbf{s}(t_*)\| > \mathcal{S} \quad (4)$$

subject to constraints on the false alarm error rate, P_{fa} , and missed detection error rate, P_{md} , that would entail if the problem were solved over many

¹ All vector norms in this work are Euclidean two-norms.

repetitions:

$$P_{fa} \leq \bar{P}_{fa} \quad (5)$$

$$P_{md} \leq \bar{P}_{md} \quad (6)$$

In the present context, a false alarm occurs if $\mathcal{H}_0$ is not rejected when $\|\mathbf{s}(t_*)\| > \mathcal{S}$, while a missed detection occurs when $\mathcal{H}_0$ is rejected when $\|\mathbf{s}(t_*)\| \leq \mathcal{S}$. The constraints $\bar{P}_{fa}$ and $\bar{P}_{md}$ are target values for the error rates, chosen according to context. At each measurement time, the solution should minimize the summed-squares of its estimation errors, such that the objective function to minimize is

$$J_k = \text{tr} \mathbf{P}_{k|k} \quad (7)$$

where $\mathbf{P}_{k|k}$ is the state error covariance at time t_k , given measurements up to and including t_k .

3 Problem Solution

Let $\mathbb{Y}_{1:k} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_k\}$, where here and in the sequel the time argument of a discrete sample is abbreviated to a subscript index. The joint density of the measurements, unconstrained by either hypothesis, can be written in a sequential form as

$$p(\mathbb{Y}_{1:k}) = p(\mathbf{y}_k | \mathbb{Y}_{1:k-1}) p(\mathbb{Y}_{1:k-1}) \quad (8)$$

As standard texts, such as Brown and Hwang [3], show, if the inputs to the estimator are zero-mean and Gaussian, then

$$p(\mathbf{y}_k | \mathbb{Y}_{1:k-1}) = N(\mathbf{y}_k - \mathbf{h}(\hat{\mathbf{x}}_{k|k-1}), \mathbf{H}_k \mathbf{P}_{k|k-1} \mathbf{H}_k^T + \mathbf{R}_k) \quad (9)$$

where $N(\boldsymbol{\xi} - \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$, $\mathbf{y}_k - \mathbf{h}(\hat{\mathbf{x}}_{k|k-1}) \equiv \boldsymbol{\varepsilon}_{k|k-1}$ is the k th filter innovation, $\mathbf{H}_k \mathbf{P}_{k|k-1} \mathbf{H}_k^T + \mathbf{R}_k \equiv \mathbf{W}_{k|k-1}$ is the innovations covariance, $\mathbf{P}_{k|k-1}$ is the filter's error covariance corresponding to $\hat{\mathbf{x}}_{k|k-1}$, and $\mathbf{H}_k = \partial \mathbf{h}_k / \partial \mathbf{x}_k|_{\hat{\mathbf{x}}_{k|k-1}}$.

3.1 Wald's Sequential Probability Ratio Test

Wald's SPRT uses a ratio of the conditional probability densities of the set of measurements $\mathbb{Y}_{1:k}$ under the alternative hypothesis, $\mathcal{H}_1$ and the null hypothesis, $\mathcal{H}_0$:

$$\Lambda_k = \frac{p(\mathbb{Y}_{1:k} | \mathcal{H}_1)}{p(\mathbb{Y}_{1:k} | \mathcal{H}_0)} = \frac{p(\mathbb{Y}_{1:k} | \|\mathbf{s}(t_*)\| > \mathcal{S})}{p(\mathbb{Y}_{1:k} | \|\mathbf{s}(t_*)\| \leq \mathcal{S})} \quad (10)$$

In the case that $p(\mathbb{Y}_{1:k} | \mathcal{H}_1) = p(\mathbb{Y}_{1:k} | \mathcal{H}_0) = 0$, Wald takes $\Lambda_k = 1$. In the 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. If $\Lambda_k \geq A$, then one should accept the

alternative hypothesis. Wald's explanations for the thresholds A and B are that we will accept the alternative hypothesis if it is A times more likely than the null, and accept the null hypothesis if it is $1/B$ times more likely than the alternative. Wald shows that such a procedure will terminate with probability one, and that

$$A \leq \frac{1 - P_{md}}{P_{fa}} \text{ and } B \geq \frac{P_{md}}{1 - P_{fa}} \quad (11)$$

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 the targeted false alarm and missed detection probabilities, as

$$A = \frac{1 - \bar{P}_{fa}}{\bar{P}_{md}} \quad \text{and} \quad B = \frac{\bar{P}_{fa}}{1 - \bar{P}_{md}} \quad (12)$$

Note however that these relations do *not* guarantee that the actual error rates obey $P_{fa} \leq \bar{P}_{fa}$ and $P_{md} \leq \bar{P}_{md}$, as required by Eqs. 5 and 6. They provide instead only the weaker inequalities [20] $P_{fa} + P_{md} \leq \bar{P}_{fa} + \bar{P}_{md}$, $P_{fa} \leq \bar{P}_{fa}/(1 - \bar{P}_{md})$, and $P_{md} \leq \bar{P}_{md}/(1 - \bar{P}_{fa})$. These inequalities guarantee that at most one of the error rates 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.

Suppose there exist a pair of norm-inequality constrained filters, one of which estimates the state at t_* under the constraint of $\mathcal{H}_1$ and the other under the complementary constraint of $\mathcal{H}_0$. Then the likelihood ratio in terms of these estimates becomes:

$$\begin{aligned} \Lambda_k &= \frac{N(\mathbf{y}_k - \mathbf{h}(\hat{\mathbf{x}}_{k|k-1, \mathcal{H}_1}), \mathbf{H}_k \mathbf{P}_{k|k-1, \mathcal{H}_1} \mathbf{H}_k^\top + \mathbf{R}_k)}{N(\mathbf{y}_k - \mathbf{h}(\hat{\mathbf{x}}_{k|k-1, \mathcal{H}_0}), \mathbf{H}_k \mathbf{P}_{k|k-1, \mathcal{H}_0} \mathbf{H}_k^\top + \mathbf{R}_k)} \Lambda_{k-1} \\ &= \frac{N(\boldsymbol{\varepsilon}_{k|k-1, \mathcal{H}_1}, \mathbf{W}_{k|k-1, \mathcal{H}_1})}{N(\boldsymbol{\varepsilon}_{k|k-1, \mathcal{H}_0}, \mathbf{W}_{k|k-1, \mathcal{H}_0})} \Lambda_{k-1} \end{aligned} \quad (13)$$

The next two subsections develop two alternatives for such filters.

3.2 Epoch-State Filter Bank

Reference [21] derives the norm-equality-constrained Kalman filter by appending a norm constraint on the state vector to Eq. (7). For present purposes, convert the norm equality constraint to an inequality constraint through the use of a slack variable, α , as described for example by Hull [9]. To this end, let the cost function to minimize be

$$J_k = \text{tr} \mathbf{P}_{k|k} + \lambda_k \left[\hat{\mathbf{x}}_{k|k}^\top \hat{\mathbf{x}}_{k|k} - \ell \pm \alpha^2 \right] \quad (14)$$

which can be written as

$$J_k = \text{tr} \left[\mathbf{P}_{k|k-1} - \mathbf{K}_k \mathbf{H}_k \mathbf{P}_{k|k-1} - \mathbf{P}_{k|k-1} \mathbf{H}_k^\top \mathbf{K}_k^\top + \mathbf{K}_k \mathbf{W}_{k|k-1} \mathbf{K}_k^\top \right] \\ + \lambda_k \left[\boldsymbol{\varepsilon}_{k|k-1}^\top \mathbf{K}_k^\top \mathbf{K}_k \boldsymbol{\varepsilon}_{k|k-1} + 2 \boldsymbol{\varepsilon}_{k|k-1}^\top \mathbf{K}_k^\top \hat{\mathbf{x}}_{k|k-1} + \hat{\mathbf{x}}_{k|k-1}^\top \hat{\mathbf{x}}_{k|k-1} - \ell \pm \alpha^2 \right] \quad (15)$$

where $\mathbf{K}_k$ is the Kalman gain matrix, $\ell = \hat{\mathbf{x}}^\top \hat{\mathbf{x}}$ is the boundary of the norm constraint, and the sign in front of α^2 is either positive to satisfy $\mathcal{H}_0$ or negative to satisfy $\mathcal{H}_1$.

The constraint equation along with the partial derivatives of the augmented cost function with respect to the Kalman gain and slack variable lead to a set of three equations and three unknowns. As Reference [9] shows, optimality conditions on the constraint boundary correspond to the solution $\alpha = 0$, and optimality conditions off the boundary correspond to $\lambda_k = 0$. Optimality conditions off the boundary ($\lambda_k = 0$) are given by the usual unconstrained Kalman filter solution, which we denote by $\hat{\mathbf{x}}_{k|k}$ and $\mathbf{P}_{k|k}$. The solution on the constraint boundary ($\alpha = 0$) is given by Reference [21]:

$$\hat{\mathbf{x}}_{k|k, (\|\hat{\mathbf{x}}_{k|k}\|=\sqrt{\ell})} = \frac{\sqrt{\ell}}{\|\hat{\mathbf{x}}_{k|k}\|} \hat{\mathbf{x}}_{k|k} \quad (16)$$

$$\mathbf{P}_{k|k, (\|\hat{\mathbf{x}}_{k|k}\|=\sqrt{\ell})} = \mathbf{P}_{k|k} + \frac{1}{\boldsymbol{\varepsilon}_{k|k-1}^\top \mathbf{W}_{k|k-1}^{-1} \boldsymbol{\varepsilon}_{k|k-1}} \left(1 - \frac{\sqrt{\ell}}{\|\hat{\mathbf{x}}_{k|k}\|} \right)^2 \hat{\mathbf{x}}_{k|k} \hat{\mathbf{x}}_{k|k}^\top \quad (17)$$

Reference [21] also shows that if the constraint applies to only a subset of the state elements, then the optimal solution for the remaining elements is just the unconstrained Kalman filter solution for those elements.

For the problem at hand, the constraint applies only at the time t_* , so one may use an epoch-state filter that sequentially processes the set of currently available measurements. In an epoch-state filter, the (unconstrained) update at t_k , for the state at a future time t_* , given a measurement innovation at t_k that is conditioned on measurements up through t_{k-1} , has the form

$$\hat{\mathbf{x}}_{*|k} = \hat{\mathbf{x}}_{*|k-1} + \mathbf{K}_k \boldsymbol{\varepsilon}_{k|k-1} \quad (18)$$

Reference [13] shows that the optimal gain for this filter is

$$\mathbf{K}_k = (\mathbf{P}_{*|k-1} - \mathbf{Q}_{*,k}) \widetilde{\mathbf{H}}_{k,*}^\top \left[\widetilde{\mathbf{H}}_{k,*} (\mathbf{P}_{*|k-1} - \mathbf{Q}_{*,k}) \widetilde{\mathbf{H}}_{k,*}^\top + \mathbf{R}_k \right]^{-1} \\ = (\mathbf{P}_{*|k-1} - \mathbf{Q}_{*,k}) \widetilde{\mathbf{H}}_{k,*}^\top \widetilde{\mathbf{W}}_{*|k-1}^{-1} \quad (19)$$

and the error covariance at epoch is given by

$$\mathbf{P}_{*|k} = \mathbf{P}_{*|k-1} - \mathbf{K}_k \widetilde{\mathbf{H}}_{k,*} (\mathbf{P}_{*|k-1} - \mathbf{Q}_{*,k}) \quad (20)$$

where $\widetilde{\mathbf{H}}_{k,*} = \mathbf{H}_k \boldsymbol{\Phi}(t_k, t_*)$, $\mathbf{Q}_{*,k}$ is the process covariance resulting from integrating the noise density, $\mathbf{Q}(t)$, *forward* from t_k to the epoch time t_* , i.e.,

$$\mathbf{Q}_{*,k} = \int_{t_k}^{t_*} \boldsymbol{\Phi}(t_*, \tau) \mathbf{Q}(\tau) \boldsymbol{\Phi}^\top(t_*, \tau) d\tau \quad (21)$$

and $\boldsymbol{\Phi}(t, t_i)$ is the solution of $\dot{\boldsymbol{\Phi}}(t, t_i) = \mathbf{A}(t) \boldsymbol{\Phi}(t, t_i)$ with $\boldsymbol{\Phi}(t_i, t_i) = \mathbf{I}$.

Application of the constraint to the initial *a priori* state and covariance yields:

$$\hat{\mathbf{s}}_{*|0, \mathcal{H}_0} = \begin{cases} \hat{\mathbf{s}}_{*|0} & \text{if } \|\hat{\mathbf{s}}_{*|0}\| \leq \mathcal{S} \\ \frac{\mathcal{S}}{\|\hat{\mathbf{s}}_{*|0}\|} \hat{\mathbf{s}}_{*|0} & \text{otherwise} \end{cases} \quad (22)$$

$$\hat{\mathbf{s}}_{*|0, \mathcal{H}_1} = \begin{cases} \hat{\mathbf{s}}_{*|0} & \text{if } \|\hat{\mathbf{s}}_{*|0}\| > \mathcal{S} \\ \frac{\mathcal{S}}{\|\hat{\mathbf{s}}_{*|0}\|} \hat{\mathbf{s}}_{*|0} & \text{otherwise} \end{cases} \quad (23)$$

$$\mathbf{P}_{ss_{*|0}, \mathcal{H}_0} = \begin{cases} \mathbf{P}_{ss_{*|0}} & \text{if } \|\hat{\mathbf{s}}_{*|0}\| \leq \mathcal{S} \\ \mathbf{P}_{ss_{*|0}} + \left(1 - \frac{\mathcal{S}}{\|\hat{\mathbf{s}}_{*|0}\|}\right)^2 \hat{\mathbf{s}}_{*|0} \hat{\mathbf{s}}_{*|0}^\top & \text{otherwise} \end{cases} \quad (24)$$

$$\mathbf{P}_{ss_{*|0}, \mathcal{H}_1} = \begin{cases} \mathbf{P}_{ss_{*|0}} & \text{if } \|\hat{\mathbf{s}}_{*|0}\| > \mathcal{S} \\ \mathbf{P}_{ss_{*|0}} + \left(1 - \frac{\mathcal{S}}{\|\hat{\mathbf{s}}_{*|0}\|}\right)^2 \hat{\mathbf{s}}_{*|0} \hat{\mathbf{s}}_{*|0}^\top & \text{otherwise} \end{cases} \quad (25)$$

The remaining elements of the state and covariance are unaffected by the constraint application.

The k th measurement update, under each hypothesis, $\mathcal{H}_i$, is as follows:

$$\hat{\mathbf{x}}_{*|k, \mathcal{H}_i} = \hat{\mathbf{x}}_{*|k-1, \mathcal{H}_i} + \mathbf{K}_{k|\mathcal{H}_i} \boldsymbol{\varepsilon}_{k|k-1, \mathcal{H}_i} \quad (26)$$

$$\mathbf{P}_{*|k, \mathcal{H}_i} = \mathbf{P}_{*|k-1, \mathcal{H}_i} - \mathbf{K}_{k|\mathcal{H}_i} \widetilde{\mathbf{H}}_{k,*} (\mathbf{P}_{*|k-1, \mathcal{H}_i} - \mathbf{Q}_{*,k}) \quad (27)$$

$$\begin{aligned} \mathbf{K}_{k|\mathcal{H}_i} &= (\mathbf{P}_{*|k-1, \mathcal{H}_i} - \mathbf{Q}_{*,k}) \widetilde{\mathbf{H}}_{k,*}^\top \left[\widetilde{\mathbf{H}}_{k,*} (\mathbf{P}_{*|k-1, \mathcal{H}_i} - \mathbf{Q}_{*,k}) \widetilde{\mathbf{H}}_{k,*}^\top + \mathbf{R}_k \right]^{-1} \\ &= (\mathbf{P}_{*|k-1, \mathcal{H}_i} - \mathbf{Q}_{*,k}) \widetilde{\mathbf{H}}_{k,*}^\top \widetilde{\mathbf{W}}_{k|k-1, \mathcal{H}_i}^{-1} \end{aligned} \quad (28)$$

Application of the constraint to the *a posteriori* state and covariance yields:

$$\hat{\mathbf{s}}_{*|k, \mathcal{H}_0} \leftarrow \begin{cases} \hat{\mathbf{s}}_{*|k, \mathcal{H}_0} & \text{if } \|\hat{\mathbf{s}}_{*|k, \mathcal{H}_0}\| \leq \mathcal{S} \\ \frac{\mathcal{S}}{\|\hat{\mathbf{s}}_{*|k, \mathcal{H}_0}\|} \hat{\mathbf{s}}_{*|k, \mathcal{H}_0} & \text{otherwise} \end{cases} \quad (29)$$

$$\hat{\mathbf{s}}_{*|k, \mathcal{H}_1} \leftarrow \begin{cases} \hat{\mathbf{s}}_{*|k, \mathcal{H}_1} & \text{if } \|\hat{\mathbf{s}}_{*|k, \mathcal{H}_1}\| > \mathcal{S} \\ \frac{\mathcal{S}}{\|\hat{\mathbf{s}}_{*|k, \mathcal{H}_1}\|} \hat{\mathbf{s}}_{*|k, \mathcal{H}_1} & \text{otherwise} \end{cases} \quad (30)$$

$$P_{ss_*|k, \mathcal{H}_0} \leftarrow \begin{cases} P_{ss_*|k, \mathcal{H}_0} & \text{if } \|\hat{s}_{*|k, \mathcal{H}_0}\| \leq \mathcal{S} \\ P_{ss_*|k, \mathcal{H}_0} + \frac{1}{\tilde{\epsilon}_{k|k-1, \mathcal{H}_0}} \left(1 - \frac{\mathcal{S}}{\|\hat{s}_{*|k, \mathcal{H}_0}\|}\right)^2 \hat{s}_{*|k, \mathcal{H}_0} \hat{s}_{*|k, \mathcal{H}_0}^\top & \text{otherwise} \end{cases} \quad (31)$$

$$P_{ss_*|k, \mathcal{H}_1} \leftarrow \begin{cases} P_{ss_*|k, \mathcal{H}_1} & \text{if } \|\hat{s}_{*|k, \mathcal{H}_1}\| > \mathcal{S} \\ P_{ss_*|k, \mathcal{H}_1} + \frac{1}{\tilde{\epsilon}_{k|k-1, \mathcal{H}_1}} \left(1 - \frac{\mathcal{S}}{\|\hat{s}_{*|k, \mathcal{H}_1}\|}\right)^2 \hat{s}_{*|k, \mathcal{H}_1} \hat{s}_{*|k, \mathcal{H}_1}^\top & \text{otherwise} \end{cases} \quad (32)$$

where

$$\tilde{\epsilon}_{k|k-1, \mathcal{H}_i} = \epsilon_{k|k-1, \mathcal{H}_i}^\top \tilde{\mathbf{W}}_{k|k-1, \mathcal{H}_1}^{-1} \epsilon_{k|k-1, \mathcal{H}_i} \quad (33)$$

The remaining elements of the state and covariance are unaffected by the constraint application.

Finally, the likelihood ratio updates sequentially as follows:

$$\Lambda_k = \frac{N(\epsilon_{k|k-1, \mathcal{H}_1}, \tilde{\mathbf{W}}_{k|k-1, \mathcal{H}_1})}{N(\epsilon_{k|k-1, \mathcal{H}_0}, \tilde{\mathbf{W}}_{k|k-1, \mathcal{H}_0})} \Lambda_{k-1} \quad (34)$$

with $\Lambda_0 = 1$.

3.3 Current-State Filter Bank

To adapt the solution described above to a current-state filtering paradigm, adopt the approach of Julier and LaViola [10], which applies a projection to the unconstrained mean and covariance using a sigma-point transformation. First, perturb each current-state constrained filter's state estimate, $\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}$, using the columns of the Cholesky factorization of its error covariance to generate a set of $2n + 1$ sigma points, $\mathcal{X}_{k|\mathcal{H}_i}$:

$$\mathcal{X}_{k|\mathcal{H}_i} = \left[\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}, \hat{\mathbf{x}}_{k|k, \mathcal{H}_i} + h\sqrt[3]{P}_{(:,1)}, \hat{\mathbf{x}}_{k|k, \mathcal{H}_i} + h\sqrt[3]{P}_{(:,2)}, \dots, \hat{\mathbf{x}}_{k|k, \mathcal{H}_i} + h\sqrt[3]{P}_{(:,n)}, \right. \\ \left. \hat{\mathbf{x}}_{k|k, \mathcal{H}_i} - h\sqrt[3]{P}_{(:,1)}, \hat{\mathbf{x}}_{k|k, \mathcal{H}_i} - h\sqrt[3]{P}_{(:,2)}, \dots, \hat{\mathbf{x}}_{k|k, \mathcal{H}_i} - h\sqrt[3]{P}_{(:,n)} \right] \quad (35)$$

where $P_{(:,i)}$ denotes column i of $\mathbf{P}_{k|k, \mathcal{H}_i}$, $\mathbf{P} = \sqrt[3]{P}\sqrt[3]{P}^\top$ denotes a Cholesky factorization, and h is the step size of the perturbations chosen for the sigma-point transformation, typically set to the square root of the kurtosis of the distribution [16, 15]. In the present case, the (inequality) constraint function is given by

$$c(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}) = \|\phi_*^s(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i})\| \leq \mathcal{S} \quad (36)$$

where the direction of the inequality is chosen according to which of the $\mathcal{H}_i$ is being enforced, and $\phi_*^s(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i})$ is the prediction of the constrained components of $\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}$ to t_* .

Reference [10] concerns equality-constrained estimation, whereas the present problem concerns inequality-constraints. As above, if $c(\mathcal{X}_{k|\mathcal{H}_0}^{(j)}) \geq \mathcal{S}$ for any sigma point arising from the null hypothesis filter, then the sigma point is projected onto the constraint boundary, and vice versa for the alternative hypothesis filter. Julier and LaViola's method for constraining the estimate allows for any projection function, $\mathbf{p}(\mathbf{x})$, that satisfies the constraint, i.e.

$$c(\mathbf{p}(\mathbf{x})) = \mathcal{S} \quad (37)$$

Therefore evaluate each sigma point, $\mathcal{X}_{k|\mathcal{H}_i}^{(j)}$, compare it to the constraint, and apply the projector:

$$\mathcal{P}_{k|\mathcal{H}_i}^{(j)} = \begin{cases} \mathbf{p}(\mathcal{X}_{k|\mathcal{H}_i}^{(j)}) & \text{if } c(\mathcal{X}_{k|\mathcal{H}_i}^{(j)}) \in \mathcal{H}_i \\ \mathcal{X}_{k|\mathcal{H}_i}^{(j)} & \text{otherwise} \end{cases} \quad (38)$$

Applying the projector, $\mathbf{p}(\mathbf{x})$, is not the innovation of Julier and LaViola's method, but rather only the first step of a two-step procedure². They argue the need for a second step by pointing out that projection does not necessarily force the mean of the resulting constrained density to satisfy the constraint. Since the mean is the property of the distribution estimated by a filter, Reference [10] suggests that the entire distribution should be translated so that the mean will satisfy the constraint. This translation, which also affects the covariance, constitutes the second step of Julier and LaViola's approach. In the context of an inequality constraint, one would only perform this step if the mean fails to satisfy the inequality. It is not obvious that the second step of Julier and LaViola's method is necessary or even desirable in the present context. By construction, the projected sigma points, $\mathcal{P}_{k|\mathcal{H}_i}^{(j)}$, will already satisfy the inequality constraint and hence are the best available representation of the distribution. Translating them in order for the mean to satisfy the constraint could therefore degrade the overall quality of the representation of the distribution.

To see that this approach is optimal, note that the minimum mean-square error (MMSE) of the state estimate that has been inequality-constrained by the projector is

$$\begin{aligned} \text{MMSE} &= \arg \min_{\mathbf{x}' \in \mathcal{H}_i} \mathbb{E} [\|\mathbf{x}' - \mathbf{x}\|^2] = \arg \min_{\mathbf{x}' \in \mathcal{H}_i} \mathbb{E} [\|\mathbf{x}' - \hat{\mathbf{x}} + \hat{\mathbf{x}} - \mathbf{x}\|^2] \\ &= \arg \min_{\mathbf{x}' \in \mathcal{H}_i} \left\{ \mathbb{E} [\|\mathbf{x}' - \hat{\mathbf{x}}\|^2] + \mathbb{E} [\|\hat{\mathbf{x}} - \mathbf{x}\|^2] - 2 \mathbb{E} [(\mathbf{x}' - \hat{\mathbf{x}})^\top (\hat{\mathbf{x}} - \mathbf{x})] \right\} \end{aligned} \quad (39)$$

and since $\hat{\mathbf{x}} = \mathbb{E} [\mathbf{x} | \mathbb{Y}_{1:k}]$, then

$$\text{MMSE} = \arg \min_{\mathbf{x}' \in \mathcal{H}_i} \mathbb{E} [\|\mathbf{x}' - \hat{\mathbf{x}}\|^2] \quad (40)$$

² The nonlinear projection method was suggested in Reference [18], but Julier and LaViola appear to be the first to suggest applying the projector using a sigma-point transformation.

Hence, the MMSE is achieved by letting $\mathbf{x}'$ be as close as possible to $\hat{\mathbf{x}}$ in the sense of the chosen two-norm, subject to the constraint that $\mathbf{x}' \in \mathcal{H}_i$. If $\hat{\mathbf{x}} \in \mathcal{H}_i$, then set $\mathbf{x}' = \hat{\mathbf{x}}$. Since the constraint is a boundary for $\mathcal{H}_i$, then if $\hat{\mathbf{x}} \notin \mathcal{H}_i$, choose the point on the constraint closest to $\hat{\mathbf{x}}$ in the sense of the chosen norm.

To conduct the WSPRT, after each measurement update, the constrained filters apply the constraint as described above to the posterior state estimate. Then merge the constrained sigma points to form the “step-1” constrained mean:

$$\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}^{(1)} = \frac{h^2 - n}{h^2} \mathcal{P}_{k|k, \mathcal{H}_i}^{(1)} + \frac{1}{2h^2} \sum_{j=2}^{2n+1} \mathcal{P}_{k|k, \mathcal{H}_i}^{(j)} \quad (41)$$

To form the associated covariance, use the divided-difference matrices of Nørgaard, et al. [16, 15], each column j of which is given by

$$\tilde{D}_{\Delta x}^{(1)}(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i})_j = \frac{1}{2h} [\mathcal{P}_{k|k, \mathcal{H}_i}^{(j+1)} - \mathcal{P}_{k|k, \mathcal{H}_i}^{(j+1+n)}] \quad (42)$$

$$\tilde{D}_{\Delta x}^{(2)}(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i})_j = \frac{\sqrt{h^2 - 1}}{2h^2} [\mathcal{P}_{k|k, \mathcal{H}_i}^{(j+1)} + \mathcal{P}_{k|k, \mathcal{H}_i}^{(j+1+n)} - 2\mathcal{P}_{k|k, \mathcal{H}_i}^{(1)}] \quad (43)$$

The application of the nonlinear constraint to the posterior covariance is then approximated (for step 1) by

$$\mathbf{P}_{k|k, \mathcal{H}_i}^{(1)} = [\tilde{D}_{\Delta x}^{(1)}(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}), \tilde{D}_{\Delta x}^{(2)}(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i})] [\tilde{D}_{\Delta x}^{(1)}(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}), \tilde{D}_{\Delta x}^{(2)}(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i})]^\top \quad (44)$$

Finally, if performing Julier and LaViola’s step 2, apply the projector to the step-1 mean and adjust its covariance, if necessary:

$$\text{if } c(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}^{(1)}) \in \mathcal{H}_i: \begin{cases} \hat{\mathbf{x}}_{k|k, \mathcal{H}_i}^{(2)} = \mathbf{p}(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}^{(1)}) \\ \mathbf{P}_{k|k, \mathcal{H}_i}^{(2)} = \mathbf{P}_{k|k, \mathcal{H}_i}^{(1)} + \left(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}^{(2)} - \hat{\mathbf{x}}_{k|k, \mathcal{H}_i}^{(1)} \right) \left(\hat{\mathbf{x}}_{k|k, \mathcal{H}_i}^{(2)} - \hat{\mathbf{x}}_{k|k, \mathcal{H}_i}^{(1)} \right)^\top \end{cases} \quad (45)$$

$$\text{otherwise:} \begin{cases} \hat{\mathbf{x}}_{k|k, \mathcal{H}_i}^{(2)} = \hat{\mathbf{x}}_{k|k, \mathcal{H}_i}^{(1)} \\ \mathbf{P}_{k|k, \mathcal{H}_i}^{(2)} = \mathbf{P}_{k|k, \mathcal{H}_i}^{(1)} \end{cases} \quad (46)$$

Finally, as before, iteratively compute the likelihood ratio for the WSPRT from the innovations of the constrained filters using Eq. (34) and compare Λ_k to Wald’s limits A and B , computed using desired missed detection and false alarm rates.

4 Discussion

As noted above, the WSPRT compares fixed hypotheses. If the state errors can randomly fluctuate, such hypotheses are no longer fixed, and other methods, such as the Shiryaev SPRT [17] may be needed. Hence, although the

solutions described above explicitly accommodate any process noise present in the system, this process noise should not enter into the states that define the hypotheses; for example, the present solutions would permit process noise that affects measurement biases, but if the constraint concerns the norm of position, then process noise that drives the true position or velocity directly would in principle invalidate the assumption of fixed hypotheses. In many real-world problems, this is not a concern, since inclusion of the process noise is a fictitious mechanism introduced to account for model uncertainties, and does not actually enter the true states.

In most successful applications of Kalman filtering to real-world nonlinear problems, some form of measurement editing is required [4]. It is primarily for this reason that inclusion of a third, unconstrained filter into the bank is suggested, as Fig. 1 indicates. The unconstrained filter can also provide a linearization reference for the two constrained filters, permitting them to operate as linearized, rather than extended Kalman filters, which could further improve consistency.

Although the epoch-state solution is conceptually simpler than the current-state solution, it has some drawbacks in terms of implementation. The most obvious is that for many applications, a current-state filter already exists, and it will be easier to retain it and “tack on” the sigma-point constraint mapping. A more fundamental potential shortcoming of the epoch-state formulation is that it uses a linear approximation, $\tilde{\mathbf{H}}_{k,*} = \mathbf{H}_k \Phi(t_k, t_*)$, to map the measurement innovations from t_k to t_* . A possible improvement, not considered in the present work, is to replace this mapping with a sigma-point transform. Also, unless there is an exact solution for the process noise integral, Eq. (21), an approximation that is suitable over longer time spans than are typical for current-state Kalman filters may be needed. Here again, sigma-point methods might be profitably employed.

Finally, one should bear in mind that even in the benign case where all models are linear, and all distributions are Gaussian, application of a constraint will introduce difficulties, which might be significant. As Reference [21] puts it,

When two random variables are related through a nonlinear transformation, it is generally impossible to exclusively relate their second moments[,] but all the moments of the original variable will contribute in the second moment of the transformed variable. Therefore, the correction on the covariance can be accurate or not, depending on the distribution of the estimation error. The [constrained covariance matrix expression] is an approximation and, like any approximation, might not be satisfactory under certain circumstances.

5 Example Application: Spacecraft Conjunctions

Let $\{\mathbf{R}_i, \mathbf{V}_i\}, i = 1, 2$ represent the inertial position and velocity of two space objects. Then, the relative position and velocity between the objects in a frame

fixed to the orbital velocity and orbit normal vectors of object 1 are

$$\mathbf{r} = \mathbf{R}_2 - \mathbf{R}_1 \quad \text{and} \quad \mathbf{v} = \mathbf{V}_2 - \mathbf{V}_1 - \boldsymbol{\omega}_1 \times \mathbf{r} \quad (47)$$

where $\boldsymbol{\omega}_1$ is the instantaneous orbital angular velocity of object 1. During a time interval of interest, the relative position may have one or more local minima, called “close approaches,” that occur when

$$\mathbf{r}^\top \mathbf{v} = 0 \quad \text{and} \quad \mathbf{r}^\top \mathbf{a} + \mathbf{v}^\top \mathbf{v} > 0 \quad (48)$$

where $\mathbf{a}$ is the difference between the total inertial accelerations of the objects. Such minima are predicted as part of a conjunction screening process, and any minima for which some additional screening criteria are met, such as $\|\mathbf{r}\|$ less than some specification, constitute events of interest for further assessment, and if at least one of the objects can maneuver, consideration of risk mitigation maneuvers. It is common to denote as t_{ca} the time at which Eqs. 48, along with any additional screening criteria, are satisfied. For purposes of applying a WSRPT to this problem, take as the null hypothesis that the miss distance is inside a spherical region bounding the maximum dimension of the two objects, known as the combined hard body sphere, at the predicted time of closest approach, and take the alternative hypothesis as the complement of the null.

Typically, one of the space objects is an active spacecraft, with a relatively well-known predictive ephemeris, while the other is a debris object, whose ephemeris is less well-known. In present practice, human operators in terrestrial mission operations centers perform nearly all conjunction assessment and consideration of risk mitigation activities. Opportunities to track and communicate with the active spacecraft will typically occur a few times per day, while debris objects may be tracked much less often, unless they pose a relatively imminent threat to an active spacecraft. Furthermore, interfaces between organizations involved in such operations typically do not provide for the sharing of tracking data. For such cases, the WSPRT of Reference [6] is more suitable than the methods of the present work. For “intramural” conjunctions within a formation or constellation of spacecraft, particularly those with inter-spacecraft communication links, and a tolerance or necessity for greater autonomy exists, the present work’s methods could be more suitable. In this latter context, this section presents a simplified example of two hypothetical spacecraft that are part of a formation that exchange “point solution” position measurements generated by Global Navigation Satellite Systems receivers. Each spacecraft redundantly processes these measurements in an unconstrained extended Kalman filter. Nominally the spacecraft maintain safe separations, but in the course of frequent formation-keeping maneuvers, execution errors may build up, and the potential for a future unsafe conjunction may develop.

5.1 Epoch-State Filter Bank Demonstration

The epoch-state filter bank demonstration of the WSPRT considers the formation in a highly elliptical orbit with period of 66.8 hours, with a close approach

of 0.5 km occurring about 30 minutes after perigee. To create the simulation truth for the demonstration, the formation was back-propagated from the close approach for two orbital revolutions using the Gim-Alfriend state transition matrix [8], which accounts for both high eccentricity and Earth oblateness. To simulate the navigation process, these measurements were corrupted by white noise with standard deviation of 10 m, and by a first-order Gauss-Markov bias process with steady-state standard deviation of 3 m. The measurements occurred at 60 sec rate. Reference [7] provides additional details regarding the simulation.

The unconstrained filter processed these measurements for just over one orbital period to achieve a converged navigation state. This filter estimated both the relative position/velocity state, and the three components of the Gauss-Markov measurement bias. Then, just after the perigee preceding the close approach, the epoch-state filter bank was initiated. The epoch-state filters operated as linearized Kalman filters, linearized about an ephemeris propagated from the unconstrained filter state that initialized them. All the filters used the simulation truth values of the noise statistics. The only process noise entering the epoch-state filters was that from the Gauss-Markov measurement bias process.

Figures 2 and 3 show the results for two cases. In each figure, the upper left subplot shows a close up of the relative motion over approximately an hour preceding the closest approach. This motion is projected into a coordinate system centered on one of the spacecraft, whose “radial” basis vector is aligned with the radius vector of that spacecraft, “crosstrack” basis vector is aligned with the orbital angular momentum vector of that spacecraft, and whose “intrack” basis vector completes a right-handed system.

In the first case, labeled “near hit,” the constraint boundary is set at 0.75 km, which is somewhat greater than the true miss distance, so an alarm should be triggered. In the second case, labeled “near miss,” the constraint boundary is set at 0.33 km, which is somewhat less than the true miss distance, so a dismissal should be triggered.

Each figure marks the epoch-state position solutions from the null hypothesis filter using red, and from the alternative hypothesis filter in green. Depending on the case, either the green or red solution resides on the constraint boundary (shown as a grey sphere), while the other solutions cluster near the true relative position, which is either inside or outside the boundary. The upper-right and lower-left subplots show the estimation errors and their one-sigma formal errors for each of the two constrained filters, with the former showing the filter that assumes the null hypothesis, and the latter showing the filter that assumes the alternative hypothesis. The lower right subplot depicts the log-likelihood ratio, and the decision bounds. In both cases, the epoch-state filters quickly reach a steady-state, dominated by the extrapolation of the measurement bias process noise, occasionally perturbed when a constraint application produces a large correction.

In Figure 2, a bias is evident in the alternative hypothesis filter, as a consequence of its solution being forced onto the constraint boundary. The null

hypothesis filter performance is less biased, and the superior performance of the null hypothesis filter leads to a rapid and definite alarm threshold crossing. In complementary fashion, Figure 3 shows a bias in the null hypothesis filter, as a consequence of its solution being forced onto the constraint boundary. The alternative hypothesis filter performance is less biased, and the superior performance of the alternative hypothesis filter leads to a rapid and definite dismissal threshold crossing.




Fig. 2 Near hit test case using constraint boundary radius of 0.75 km when actual miss distance is 0.5 km.




Fig. 3 Near miss test case using constraint boundary radius of 0.33 km when actual miss distance is 0.5 km.

5.2 Current-State Filter Bank Demonstration

Application of the current-state filter bank requires the specification of a projector. For the present case, we choose a projector by solving Lambert's problem for the velocity of one spacecraft at time t_k that will adjust the length of the relative position at $t_* = t_{ca}$ so that it lies on the constraint boundary, as necessary to satisfy the inequality constraint. For two-body orbital motion, the solution may be expressed using classical f and g functions formulated using universal variables [2]. For specificity, this application perturbs the velocity of spacecraft 1, although either spacecraft's velocity could be perturbed. The

velocity perturbation is

$$\mathbf{V}_1(t_k) = \frac{\mathbf{R}_1^*(t_*) - f\mathbf{R}_1(t_k)}{g} \quad (49)$$

where $\mathbf{R}_1^*(t_*)$ is the targeted position of one spacecraft such that the relative position vector at t_* is lengthened or shortened to match the constraint, $\mathcal{S}$, which is the combined hard body radius:

$$\mathbf{R}_1^*(t_*) = \mathbf{R}_2(t_*) - \mathcal{S} \frac{\mathbf{r}_*}{\|\mathbf{r}_*\|} \quad (50)$$

Now specifying the contents of the state vector as $\mathbf{x} = [\mathbf{R}_1^\top, \mathbf{V}_1^\top, \mathbf{R}_2^\top, \mathbf{V}_2^\top]^\top$, the projector is given by

$$\mathbf{p}(\mathbf{x}_k) = \begin{bmatrix} \mathbf{R}_1(t_k) \\ \frac{\mathbf{R}_1^*(t_*) - f\mathbf{R}_1(t_k)}{g} \\ \mathbf{R}_2(t_k) \\ \mathbf{V}_2(t_k) \end{bmatrix} \quad (51)$$

for those $\mathbf{x}_k$ that do not satisfy the appropriate inequality constraint for the current hypothesis. Otherwise, the projector leaves states that already satisfy the constraint alone.

The current-state filter bank demonstration again considers the formation in a highly elliptical orbit, this time with period of about 24 hours, with a close approach about 30 minutes before perigee. The combined hard-body radius constraint is fixed at 120 m. The current-state filters are standard extended Kalman filters, modeling the full nonlinear two-body dynamics, including a small, fictitious acceleration process noise, with a constant isotropic intensity of 1×10^{-9} m/s^{5/2} per axis. Each spacecraft's point solution measurements are bias-free, corrupted only by isotropic white noise of 1 m per component. Reference [5] provides additional details regarding the simulation.

Similar to the previous subsection, Figures 4 and 5 present near-hit and near-miss cases, with similar plotting conventions, except as noted. In these plots, dashed lines show the filters' 3σ formal error envelope. For reference, gray solid and dashed lines show the unconstrained filter's estimate and formal error envelope. Once the WSPRT produces a recommendation, the filters stop taking measurements and simply propagate to t_{ca} . The plots show that the fictitious process noise these filters utilize drives the performance of their predictions to t_{ca} .

In Figure 4, the near hit case, the null hypothesis filter closely tracks the unconstrained filter, leading to a rapid and definitive conclusion to raise the alarm and recommend a conjunction avoidance maneuver. In Figure 5, the near miss case, neither of the constrained filters predicts as accurately as the unconstrained filter, but the alternative hypothesis filter is slightly more consistent with its own formal errors than the null hypothesis filter, leading to a dismissal.

6 Conclusions

This work has presented two optimizing solutions to a decision problem. The decision concerns choosing between fixed, compound hypotheses that divide a state space according to a two-norm constraint. Both proposed solutions involve novel versions of Wald's sequential probability ratio test that use Kalman filter banks, whose members are inequality constrained by a two-norm. The first solution uses epoch-state filters, while the second solution uses current-state filters. The first solution applies the constraint to the epoch state filters using a slack variable, while the second solution applies the constraint to the current-state filters using a sigma-point projector.

While the epoch-state approach is conceptually simpler, the current-state approach might be easier in practice to implement, depending the application. The current-state approach also contains fewer linearizations, potentially widening its applicability.

This problem was motivated by the topic of space object conjunction assessment, where the norm constraint applies to the distance between two spacecraft and/or orbital debris objects at their point of closest approach. Demonstrations using simplified spacecraft conjunction examples showed successful performance by both solution approaches.

7 Acknowledgement

Portions of this work were presented as References [7] and [5].

References

1. Akella, M.R., Alfriend, K.T.: Probability of collision between space objects. *Journal of Guidance, Control, and Dynamics* **23**(5), 769–772 (2000). DOI 10.2514/2.4611
2. Battin, R.H.: An Introduction to the Mathematics and Methods of Astrodynamics, 1st edn. AIAA Education Series. AIAA, New York (1987). ISBN: 0930403258
3. Brown, R.G., Hwang, P.Y.: Introduction to Random Signals and Applied Kalman Filtering, 3rd edn. John Wiley and Sons, Inc., New York, NY (1997). ISBN: 0471128392
4. Carpenter, J.R., D'Souza, C.N. (eds.): Navigation Filter Best Practices, 1st edn. No. NASA/TP–2018–219822 in NASA Technical Publications (<https://ntrs.nasa.gov>). National Aeronautics and Space Administration, Washington, DC 20546-0001 (2018)
5. Carpenter, J.R., Markley, F.L.: Current-state constrained filter bank for Wald testing of spacecraft conjunctions. In: T. Martin-Mur (ed.) Proceedings of the 23rd International Symposium on Space Flight Dynamics. www.issfd.org (2012)
6. Carpenter, J.R., Markley, F.L.: Wald sequential probability ratio test for space object conjunction assessment. *Journal of Guidance, Control, and Dynamics* **37**(5), 1385–1396 (2014). DOI 10.2514/1.G000478
7. Carpenter, J.R., Markley, F.L., Alfriend, K.T., Wright, C., Arcido, J.: Sequential probability ratio test for collision avoidance maneuver decisions based on a bank of norm-inequality-constrained epoch-state filters. In: Astrodynamics 2011, *Advances in the Astronautical Sciences*, vol. 142, pp. 1105–1124. Univelt (2012). ISBN: 978-0-87703-577-0
8. Gim, D.W., Alfriend, K.T.: State transition matrix of relative motion for the perturbed noncircular reference orbit. *Journal of Guidance, Control, and Dynamics* **26**(6), 956–971 (2003). DOI 10.2514/2.6924

9. Hull, D.G.: Optimal Control Theory for Applications. Springer (2003). Isbn: 0387400702
10. Julier, S.J., LaViola, J.J.: On Kalman filtering with nonlinear equality constraints. *IEEE Transactions on Signal Processing* **55**(6), 2774–2784 (2007). DOI 10.1109/TSP.2007.893949
11. Kulldorff, M., Davis, R., Kolczak, M., Lewis, N., Lieu, T., Platt, R.: A maximized sequential probability ratio test for drug and vaccine safety surveillance. *Sequential Analysis* **30** (2011). DOI 10.1080/07474946.2011.539924
12. Magill, D.: Optimal adaptive estimation of sampled stochastic processes. *IEEE Transactions on Automatic Control* **10**(4), 434–439 (1965). DOI 10.1109/TAC.1965.1098191
13. Markley, F.L., Carpenter, J.R.: Linear covariance analysis and epoch state estimators. *Journal of the Astronautical Sciences* **59**(3), 585–605 (2012). DOI 10.1007/s40295-014-0006-0
14. Montenbruck, O.: An epoch state filter for use with analytical orbit models of low earth satellites. *Aerospace Science and Technology* **4**(4), 277–287 (2000). DOI 10.1016/S1270-9638(00)00133-4
15. Nørgaard, M., Poulsen, N.K., Ravn, O.: Advances in derivative-free state estimation for nonlinear systems. Tech. Rep. IMM-REP-1998-15, Technical University of Denmark, 2800 Lyngby, Denmark (2000)
16. Nørgaard, M., Poulsen, N.K., Ravn, O.: New developments in state estimation for nonlinear estimation problems. *Automatica* **36**(11), 1627–1638 (2000). DOI 10.1016/S0005-1098(00)00089-3
17. Shiryaev, A.N.: Optimal Stopping Rules, *Applications of Mathematics*, vol. 8, 1st edn., chap. 4. Springer (1978). ISBN:3540740104
18. Simon, D., Chia, T.L.: Kalman filtering with state equality constraints. *IEEE Transactions on Aerospace and Electronic Systems* **38**(1), 128–136 (2002). DOI 10.1109/7.993234
19. Speyer, J.L., White, J.E.: Shiryaev sequential probability ratio test for redundancy management. *Journal of Guidance, Control and Dynamics* **7**(5), 588–595 (1984). DOI 10.2514/3.19899
20. Wald, A.: Sequential Analysis, chap. 1–3. Dover Publications (2004). ISBN:0486439129
21. Zanetti, R., Majji, M., Bishop, R.H., Mortari, D.: Norm-constrained Kalman filtering. *Journal of Guidance, Control, and Dynamics* **32**(5), 1458–1465 (2009). DOI 10.2514/1.43119




Fig. 4 Near hit case ($\|\mathbf{r}(t_{ca})\| = 114$ m)




Fig. 5 Near miss case ($\|\mathbf{r}(t_{ca})\| = 132$ m)