SAND2022-13122C

A Monitor for Correlated Kalman Filter Innovations

Tucker Haydon, *Sandia National Labs*
Connor Brashar, *Sandia National Labs*

Biographies

Tucker Haydon is a Senior Member of the Technical Staff at Sandia National Laboratories. He works on technologies related to inertial navigation, global navigation satellites systems, and terrain-aided navigation.

Connor Brashar has 8+ years of experience in navigation warfare, GPS signal processing, and sensor fusion. He is currently a Senior Electrical Engineer at Sandia Laboratories, where he is a technical lead for Sandia's GPS-Denied Research Imperative. He leads several projects in GPS-contested navigation technologies.

Abstract

A novel monitor is introduced to detect correlations in a Kalman filter's pre-fit measurement innovations. The two primary innovation monitors used today – innovations Snapshot and Sequence monitors – are excellent tools for detecting system faults and biases. These monitors, however, are insensitive to correlations between a Kalman filter's innovations – a key metric in evaluating and validating system assumptions and performance. A new monitor is proposed which is specifically sensitive to these correlations. The monitor evaluates the sample covariance matrix of a Kalman filter's innovations over a finite horizon and employs a hypothesis test to determine if a modeling fault has occurred. The monitor is theoretically developed and then validated with two simulations. In the first simulation, correlated random samples are drawn from a multivariate Gaussian distribution, and it is demonstrated that the proposed monitor raises a true-positive flag significantly more often than standard Snapshot and Sequence monitors while maintaining the same false-positive ratio. In the second simulation, a Monte Carlo evaluation of a simple, two-dimensional, GNSS-like example is presented wherein the presented monitor effectively detects correlation modeling errors while a Sequence monitor does not. The proposed correlation monitor has potential applications in atmospheric monitoring, navigation receiver clock monitoring, and GNSS anti-spoofing among others – essentially any application in which correlated faults can occur.

1. INTRODUCTION

The importance of accurate, trusted, and reliable navigation estimates cannot be overstated, especially for safety-of-life applications such as civilian airliner navigation, guidance, and control. For estimates to be trusted and reliable, navigation systems must confidently navigate even in the presence of potentially faulty sensor measurements and models. Navigation systems must automatically identify system failures and mitigate them before they can impact the navigation solution. The two most common techniques employed for automated fault detection are Kalman filter innovations Snapshot and Sequence monitors (Groves, 2013). These monitors are effective at detecting egregious sensor faults and biases, respectively, but are both ineffective at detecting correlated sensor errors. Less well-known and rarely employed are monitors that specifically tuned to detect correlations in a Kalman filter's innovations. Ultimately, to guarantee an accurate, trusted, and reliable navigation solution, all models, assumptions, and measurements should be scrutinized by a battery of monitors, constantly monitoring for any systematic changes to means, variances, or covariances. Here, a new correlations monitor is proposed to add to such a battery.

The concept of innovations monitoring dates to near the inception of Kalman filtering (Mehra & Peschon, 1971). These early publications established the primary methods by which the underlying model assumptions of a Kalman filter are tested. The two most popular methods – innovations Snapshot and Sequence monitors – effectively monitor the scalar variance and mean, respectively, of a filter's innovations for any suspicious deviations. These two methods have been employed across a variety of fields, subjects, and applications including GNSS measurement monitoring (Quartararo & Langel, 2021; Tanil et al., 2018)

and anti-spoofing (Liu et al., 2019; Tanil et al., 2016), among countless others. Less well-known are innovation monitors which scrutinize the vector *covariance* of a Kalman filter's innovations. These monitors accumulate the sample covariance matrix of subsequent innovations before reducing it to a scalar value needed to make a fault determination. Previous efforts have used the matric trace (Mehra & Peschon, 1971), the sum of the matrix elements (Hajiyev, 2008), and the maximum eigenvalue (Hajiyev & Caliskan, 1998) as the matrix-to-scalar reduction. Here, a different matrix-to-scalar reduction is introduced which is based on the difference between the matrix trace and its log-determinant. While this scalar reduction is not new to the multivariate statistics community (Anderson, 1958; Giri, 1977; Muirhead, 2009), it has not yet been applied to Kalman filtering innovations monitoring. Hajiyev (2008) referenced other matrix-to-scalar reduction techniques such as the one proposed here, but immediately discounted them due to asymptotic requirements. It is demonstrated, however, that these asymptotic requirements may be satisfied within a relatively short measurement horizon, enabling this new monitor to be used with mild measurement horizon requirements.

This publication makes the following three contributions. First, the ineffectiveness of Snapshot and Sequence monitors at detecting correlations in innovations is plainly demonstrated, motivating the need for correlation monitors. Second, a novel innovations monitor – the Sphericity monitor – is introduced which specifically watches for correlations in a Kalman filter's measurement innovations. Finally, the effectiveness of this monitor is demonstrated with a simple GNSS-like simulation, evidencing the capabilities of the Sphericity monitor towards detecting correlated errors in range-based measurements.

2. INNOVATION MONITORS

Let $\boldsymbol{x}_k^* \in \mathbb{R}^N$, $\overline{\boldsymbol{x}}_\boldsymbol{k} \in \mathbb{R}^N$, and $\overline{\boldsymbol{P}}_\boldsymbol{k} \in \mathbb{R}^{N\times N}$ be respectively the true state, estimated state, associated estimated error covariance of a Kalman filter at time t_k. The estimated state is modeled as a random multivariate Gaussian vector centered on the true state with covariance $\overline{\boldsymbol{P}}_\boldsymbol{k}$:

$$\overline{\boldsymbol{x}}_\boldsymbol{k} \sim \boldsymbol{\mathcal{N}}\left(\boldsymbol{x}_\boldsymbol{k}^*, \overline{\boldsymbol{P}}_\boldsymbol{k}\right) \tag{1}$$

Suppose the state were observed by a measurement $\boldsymbol{z}_\boldsymbol{k} \in \mathbb{R}^M$ modeled with the following linear (or sufficiently linearized) relationship:

$$\boldsymbol{z}_\boldsymbol{k} = \boldsymbol{H}_\boldsymbol{k}\boldsymbol{x}_\boldsymbol{k}^* + \boldsymbol{v}_\boldsymbol{k} \tag{2}$$

where $\boldsymbol{v}_\boldsymbol{k} \in \mathbb{R}^M$ is random measurement noise drawn from a zero-mean multivariate Gaussian distribution with covariance $\boldsymbol{R}_\boldsymbol{k} \in \mathbb{R}^{M\times M}$, and $\boldsymbol{H}_\boldsymbol{k} \in \mathbb{R}^{M\times N}$ is the linear observation map.

It is further assumed that the measurement noise and state estimation error are independent:

$$\boldsymbol{E}\left[\left(\boldsymbol{x}_\boldsymbol{k}^* - \overline{\boldsymbol{x}}_\boldsymbol{k}\right)\boldsymbol{v}_\boldsymbol{j}^\boldsymbol{T}\right] = \boldsymbol{0}\ \forall\ j,k \tag{3}$$

The measurement innovation $\tilde{\boldsymbol{y}}_\boldsymbol{k} \in \mathbb{R}^M$ is defined:

$$\tilde{\boldsymbol{y}}_\boldsymbol{k} = \tilde{\boldsymbol{z}}_\boldsymbol{k} - \boldsymbol{H}_\boldsymbol{k}\overline{\boldsymbol{x}}_\boldsymbol{k} \tag{4}$$

where $\tilde{\boldsymbol{z}}_\boldsymbol{k} \in \mathbb{R}^M$ is the realized measurement at t_k.

The measurement covariance $\boldsymbol{S}_\boldsymbol{k} \in \mathbb{R}^{M\times M}$ follows from the definition of the measurement innovation:

$$\boldsymbol{S}_\boldsymbol{k} = \boldsymbol{H}_\boldsymbol{k}\overline{\boldsymbol{P}}_\boldsymbol{k}\boldsymbol{H}_\boldsymbol{k}^\boldsymbol{T} + \boldsymbol{R}_\boldsymbol{k} \tag{5}$$

The normalized measurement innovation $\hat{\boldsymbol{y}}_\boldsymbol{k} \in \mathbb{R}^M$ is defined:

$$\hat{\boldsymbol{y}}_\boldsymbol{k} = \tilde{\boldsymbol{S}}_\boldsymbol{k}^{-\boldsymbol{1}}\tilde{\boldsymbol{y}}_\boldsymbol{k} \tag{6}$$

where $\tilde{\boldsymbol{S}}_\boldsymbol{k} \in \mathbb{R}^{M\times M}$ is the square root of $\boldsymbol{S}_\boldsymbol{k}$ such that:

$$\boldsymbol{S}_\boldsymbol{k} = \tilde{\boldsymbol{S}}_\boldsymbol{k}\tilde{\boldsymbol{S}}_\boldsymbol{k}^\boldsymbol{T} \tag{7}$$

It is well-known that, under the model assumptions in Equations (1) - (3), the normalized measurement innovations $\hat{\boldsymbol{y}}_\boldsymbol{k} \in \mathbb{R}^M$ are distributed like a unit, independent, uncorrelated multivariate Gaussian (see Appendix A for a proof):

$$\hat{\boldsymbol{y}}_{\boldsymbol{k}} \sim \boldsymbol{\mathcal{N}}(\boldsymbol{0}, \boldsymbol{I}) \tag{8}$$

This Gaussian distribution forms the basis of nearly all innovations monitors. From this distribution, several important properties are identified: (1) scalar innovation components should have unit variance; (2) scalar innovation components should be zero mean; and (3) all innovations should be independent, both within and between measurement epochs (see Appendix A for a proof).

The innovations Snapshot monitor is effective at testing the first property. For every scalar innovation, the Snapshot monitor forms the following test statistic $\Lambda^* \in \mathbb{R}$:

$$\Lambda^* = \hat{y}_k^{(i)} \tag{9}$$

where $\hat{y}_k^{(i)} \in \mathbb{R}$ is the ith scalar component of $\hat{\boldsymbol{y}}_{\boldsymbol{k}}$. Under the null hypothesis (H_0) that the innovations are zero-mean, unit multivariate Gaussians, this test statistic should itself be distributed as a scalar zero-mean, unit Gaussian. This test statistic and associated distribution are sufficient to create a hypothesis test to determine the likelihood of the innovation under H_0 with a prescribed false-positive ratio $\alpha \in \mathbb{R}$. Effectively, the Snapshot monitor evaluates the standard score of the innovation and raises a flag if it exceeds a threshold. While effective at detecting if the variance of an innovation is non-unit, this monitor is ineffective at detecting innovations with near-unit variance and a small bias – that is $\hat{y}_k^{(i)} \sim \mathcal{N}(\epsilon, 1)$ where ϵ is a non-zero constant.

To detect such an error (and indeed test the second innovations property), an innovations Sequence monitor is employed. A Sequence monitor accumulates scalar innovations over a finite horizon $L \in \mathbb{Z}$ and forms the following test statistic $\Lambda^* \in \mathbb{R}$:

$$\Lambda^* = \sum_{k=1}^{L} \sum_{i=1}^{i=M} \hat{y}_k^{(i)^2} \tag{10}$$

Under the null hypothesis (H_0) that the innovations are zero-mean, unit multivariate Gaussians, this test statistic should be distributed as a chi-squared distribution with LM degrees of freedom. This test statistic and associated distribution are sufficient to create a hypothesis test to determine the likelihood of the innovation under H_0 with a prescribed false-positive ratio $\alpha \in \mathbb{R}$. Effectively, the Sequence monitor evaluates the sum-square of the innovations and raises a flag if it exceeds a threshold. If the innovations are slightly biased, then the bias is square-summed until it exceeds the threshold bounds. This monitor effectively increases the observability of any biases by accumulating them out of the noise floor.

Neither of these monitors, however, is effective at testing the third innovation property – that the innovations should be independent. To address this, the innovations Sphericity monitor is introduced and subsequently evidenced. The Sphericity monitor first calculates the scaled sample covariance $\boldsymbol{B} \in \mathbb{R}^{M \times M}$ of innovations over a finite horizon $L \in \mathbb{Z}$:

$$\boxed{\boldsymbol{B} = \sum_{k=1}^{L} (\hat{\boldsymbol{y}}_{\boldsymbol{k}} - \overline{\boldsymbol{y}})(\hat{\boldsymbol{y}}_{\boldsymbol{k}} - \overline{\boldsymbol{y}})^T} \tag{11}$$

where $\overline{\boldsymbol{y}} \in \mathbb{R}^M$ is the sample mean of the innovations over the horizon. Then, the Sphericity monitor reduces the matrix $\boldsymbol{B}$ to the following scalar test statistic $\Lambda^* \in \mathbb{R}$:

$$\boxed{\Lambda^* = -LM\big(1 - \ln(L)\big) - L \ln\big(\det(\boldsymbol{B})\big) + \mathrm{tr}(\boldsymbol{B})} \tag{12}$$

Under the null hypothesis (H_0) that the innovations are zero-mean, unit multivariate Gaussians, this test statistic should be distributed as a chi-squared distribution with $M(M+1)/2$ degrees of freedom (Anderson, 1958). A proof of this distribution is given in Appendix B. Two brief notes are stated regarding the Sphericity monitor.

First, the name 'Sphericity' is given to the monitor to reflect the innovations property that it tests. If the innovations are indeed independent, then the contours of equal probability density represented by sample covariance matrix should be hyperspherical in shape. If the innovations were instead correlated, then these contours would be hyperellipsoidal in shape. Hence, the Sphericity monitor effectively evaluates whether the sample covariance ellipsoid is sufficiently spherical in shape.

Second, the Sphericity monitor requires that the accumulation horizon is large ($L \gg 1$) before the resulting distribution converges to a chi-squared distribution (see the derivation in Appendix B). Furthermore, this ensures that any detected correlations are not due to spurious correlations in random noise and are in fact due to modeling errors. Initial simulations

indicate that L need only be several multiples greater than the number of measurement channels – that is, $L \geq nM$ where n is on the order of 5 or more.

The effectiveness of this monitor at detecting correlated innovations is illustrated in the following two sections.

3. GAUSSIAN SAMPLING SIMULATION

The effectiveness of the Sphericity monitor and comparative ineffectiveness of the Snapshot and Sequence monitors at detecting correlated innovations is illustrated in the following simple Gaussian sampling simulation.

Random samples were drawn from a zero-mean, 2D multivariate Gaussian distribution with covariance

$$\Sigma = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} \qquad (13)$$

where $\rho \in [0, 1]$ is a currently unspecified correlation parameter. When $\rho = 0$, this distribution mirrors that of a normalized measurement innovation distribution under the model assumptions enumerated in Equations (1) - (3).

Innovation Snapshot, Sequence, and Sphericity monitors were evaluated on samples drawn from the distribution. The Sequence and Sphericity monitors each accumulated 100 samples before engaging their respective hypothesis tests while the Snapshot monitor engaged its hypothesis test after only a single sample. Each test was run 100,000 times in a Monte Carlo fashion for two scenarios: $\rho \in \{0.0, 0.5\}$. The first scenario establishes a true-negative baseline – no sample correlations exist and none should be detected. The second scenario establishes a true-positive baseline – significant sample correlations exist and should be detected.

Figure 1 tabulates the results of the two Monte Carlo simulations with several unraveled-and-stacked confusion matrices. A perfect monitor is listed to establish the desired baseline confusion matrix result. A perfect monitor would detect true positives and negatives 100 percent of the time (a ratio of 1.0) and never any false positives or negatives. Both the Snapshot and Sequence monitors were ineffective at detecting sample correlations – they both had a true-positive ratio of less than 2 percent. In contrast, the Sphericity monitor was highly effective at detecting sample correlations – a true-positive ratio of 99 percent!

	$\rho^* = 0.0$		$\rho^* = 0.5$	
	True Negative	False Positive	True Positive	False Negative
Perfect Monitor	1.0	0.0	1.0	0.0
Snapshot Monitor	0.99	0.01	0.01	0.99
Sequence Monitor	0.99	0.01	0.02	0.98
Sphericity Monitor	0.99	0.01	0.99	0.01

Figure 1: *Unraveled confusion matrix for the Gaussian sampling Monte Carlo. The Snapshot and Sequence monitors are ineffective at detecting sample correlations (true positives). The Sphericity monitor is highly effective at detecting sample correlations while maintaining a low false-positive ratio.*

It is not surprising that the Snapshot monitor is ineffective at detecting sampling correlations: the component variance of each of the scalar samples is unity and thus adheres to the null hypothesis of the Snapshot monitor. Another perspective – a test that evaluates the scalar components of a sample irrespective of any other components is of course insensitive to inter-sample correlations.

Perhaps more interesting, the Sequence monitor is rather ineffective at detecting sample correlations. One potential reason for this lies in the fact that the sample mean is still zero – no mean value is square-summed above the threshold bounds as in a sample bias case. Another perspective – under a zero-mean assumption, the Sequence monitor effectively sums the variances of the scalar samples and evaluates how different it is from a unit variance sum. Given that the component variances are unit themselves, its unsurprising that the hypothesis test is ineffective at correctly rejecting the null hypothesis.

In contrast, the Sphericity monitor is specifically sensitive to measurement correlations as a result of the sample covariance calculation. Unlike the Snapshot and Sequence monitors which scrutinize the samples as *scalars*, the Sphericity monitor maintains the *vector relationships* between the samples while accumulating the sample covariance. Only after the finite horizon accumulation does the Sphericity monitor then reduce the samples to a scalar test statistic. In this manner, the Sphericity monitor is effective at detecting sample correlations while the other monitors are not.

4. SIMPLE 2D SIMULATION

A simple two-dimensional (2D) simulation was created to demonstrate how the sphericity monitor might be employed in a GNSS receiver and to further evidence its capabilities. A pair of Monte Carlo simulations was run: one with measurement error correlations properly modeled and one with the measurement errors incorrectly assumed to be independent. Sequence and Sphericity innovations monitors were initialized for each of these Monte Carlo simulations and accumulated the innovations over an ever-increasing horizon. It is demonstrated that the Sphericity monitor is effective at detecting the modeling error while the Sequence monitor is not.

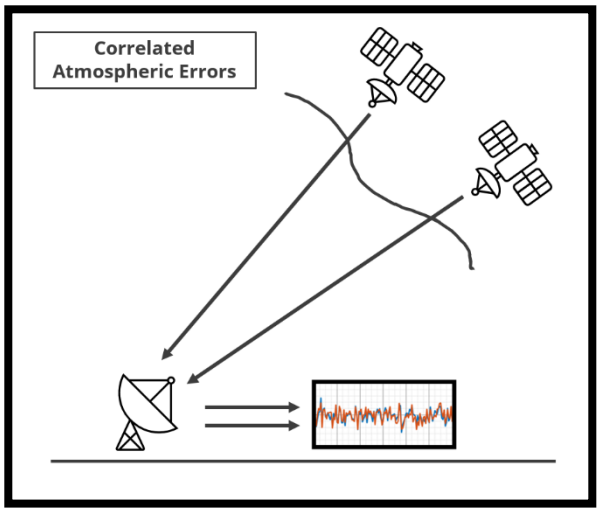

Figure 2*: A 2D receiver measures the range to two satellite vehicles. The range measurements are corrupted by localized atmospheric effects which induce correlated errors in the measurement pairs.*

Consider a simple 2D example wherein a receiver estimates its position from a pair of range measurements to two overhead stationary satellite vehicles (SV) (see Figure 2). In this example, clock biases are ignored to make the problem observable. A navigation engineer models the range measurements at time $t_k \in \mathbb{R}$ as the sum of the geometric range and additive, independent, uncorrelated Gaussian noise:

$$z_k^{(i)} = \left\| \boldsymbol{r_k^{(i)}} \right\| + \epsilon_k^{(i)} \tag{14}$$

$$\boldsymbol{R} = \boldsymbol{E}\left[\boldsymbol{\epsilon_k}\boldsymbol{\epsilon_k^T}\right] = \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix} \tag{15}$$

where $z_k^{(i)} \in \mathbb{R}$ is the range measurement to the ith SV, $\boldsymbol{r_k^{(i)}}$ is the range vector to the ith SV, $\epsilon^{(i)} \in \mathbb{R}$ is the ith component of the Gaussian noise vector $\boldsymbol{\epsilon_k} \in \mathbb{R}^2$, $\sigma^2 \in \mathbb{R}$ is the noise variance, and $\boldsymbol{R} \in \mathbb{R}^{2\times 2}$ is the modeled error covariance.

Unknown to the navigation engineer, localized atmospheric effects induce correlations in the measurement errors. The true covariance of the range measurements is:

$$\boldsymbol{R^*} = \boldsymbol{E}\left[\boldsymbol{\epsilon_k}\boldsymbol{\epsilon_k^T}\right] = \begin{bmatrix} \sigma^2 & \rho^*\sigma^2 \\ \rho^*\sigma^2 & \sigma^2 \end{bmatrix} \tag{16}$$

where $\rho^* \in [0, 1]$ is the measurement correlation coefficient and $\boldsymbol{R^*} \in \mathbb{R}^{2\times 2}$ is the true noise covariance. Here successive measurement errors are still assumed to be independent – that is, while atmospheric effects introduce similar errors between two range measurements at a given time t_k, these errors imply no information about measurement errors at some other time t_l.

An extended Kalman filter (EKF) is employed to estimate the position of the receiver. The position of the receiver, $\boldsymbol{x}(t) \in \mathbb{R}^2$, is modeled as a random walk:

$$\frac{d}{dt}\boldsymbol{x}(\boldsymbol{t}) = \boldsymbol{w}(\boldsymbol{t}) \tag{17}$$

where $\boldsymbol{w}(t) \in \mathbb{R}^2$ is random, independent white noise with a two-sided power spectral density $\boldsymbol{Q} \in \mathbb{R}^{2\times 2}$.

The dynamics are linear and straightforward to discretize over a period $\Delta t \in \mathbb{R}$. Let $\hat{\boldsymbol{x}}_{\boldsymbol{k-1}} \in \mathbb{R}^2$ and $\hat{\boldsymbol{P}}_{\boldsymbol{k-1}} \in \mathbb{R}^{2\times 2}$ be respectively the posterior state estimate and associated error covariance matrix of a Kalman filter at time $t_{k-1} \in \mathbb{R}$. Furthermore, let $\overline{\boldsymbol{x}}_{\boldsymbol{k}} \in \mathbb{R}^2$ and $\overline{\boldsymbol{P}}_{\boldsymbol{k}} \in \mathbb{R}^{2\times 2}$ be respectively the propagated prior state estimate and associated error covariance of a Kalman filter at time $t_k \in \mathbb{R}$. These estimates are related by the following linear model:

$$\overline{\boldsymbol{x}}_{\boldsymbol{k}} = \boldsymbol{F}_{\boldsymbol{d}}\hat{\boldsymbol{x}}_{\boldsymbol{k-1}} \tag{18}$$

$$\overline{\boldsymbol{P}}_{\boldsymbol{k}} = \boldsymbol{F}_{\boldsymbol{d}}\hat{\boldsymbol{P}}_{\boldsymbol{k-1}}\boldsymbol{F}_{\boldsymbol{d}}^{\boldsymbol{T}} + \boldsymbol{Q}_{\boldsymbol{d}} \tag{19}$$

where $\boldsymbol{F}_{\boldsymbol{d}} \in \mathbb{R}^{2\times 2}$ and $\boldsymbol{Q}_{\boldsymbol{d}} \in \mathbb{R}^{2\times 2}$ are the respective discretized state transition and propagation error covariance matrices:

$$\boldsymbol{F}_{\boldsymbol{d}} = \boldsymbol{I}_{\boldsymbol{2\times 2}} \tag{20}$$

$$\boldsymbol{Q}_{\boldsymbol{d}} = \boldsymbol{Q}\Delta t \tag{21}$$

The measurement model from Equation (14) is linearized to integrate the range measurements into the EKF:

$$\boldsymbol{H}_{\boldsymbol{k}} = \left[\frac{\boldsymbol{r}_{\boldsymbol{k}}^{(\boldsymbol{1})}}{\|\boldsymbol{r}_{\boldsymbol{k}}^{(\boldsymbol{1})}\|} \quad \frac{\boldsymbol{r}_{\boldsymbol{k}}^{(\boldsymbol{2})}}{\|\boldsymbol{r}_{\boldsymbol{k}}^{(\boldsymbol{2})}\|}\right]^T \tag{22}$$

All the required equations have been developed to simulate the EKF. The standard Kalman filter equations with the addition of the normalized measurement innovation calculation are listed for convenience:

Predict	**Update**
	$\tilde{\boldsymbol{y}}_{\boldsymbol{k}} = \tilde{\boldsymbol{z}}_{\boldsymbol{k}} - \boldsymbol{H}_{\boldsymbol{k}}\overline{\boldsymbol{x}}_{\boldsymbol{k}}$
	$\boldsymbol{S}_{\boldsymbol{k}} = \boldsymbol{H}_{\boldsymbol{k}}\overline{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{H}_{\boldsymbol{k}}^{\boldsymbol{T}} + \boldsymbol{R}_{\boldsymbol{k}}$
	$\tilde{\mathbf{S}}_{\mathbf{k}} = \mathrm{chol}(\boldsymbol{S}_{\boldsymbol{k}}, \mathrm{'lower'})$
$\overline{\boldsymbol{x}}_{\boldsymbol{k}} = \boldsymbol{F}_{\boldsymbol{d}}\hat{\boldsymbol{x}}_{\boldsymbol{k-1}}$	$\hat{\mathbf{y}}_{\mathbf{k}} = \tilde{\mathbf{S}}_{\mathbf{k}}^{-\mathbf{1}}\tilde{\mathbf{y}}_{\boldsymbol{k}}$
$\overline{\boldsymbol{P}}_{\boldsymbol{k}} = \boldsymbol{F}_{\boldsymbol{d}}\hat{\boldsymbol{P}}_{\boldsymbol{k-1}}\boldsymbol{F}_{\boldsymbol{d}}^{\boldsymbol{T}} + \boldsymbol{Q}_{\boldsymbol{d}}$	$\boldsymbol{K}_{\boldsymbol{k}} = \overline{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{H}_{\boldsymbol{k}}^{\boldsymbol{T}}\boldsymbol{S}_{\boldsymbol{k}}^{-\boldsymbol{1}}$
	$\hat{\boldsymbol{x}}_{\boldsymbol{k}} = \overline{\boldsymbol{x}}_{\boldsymbol{k}} + \boldsymbol{K}_{\boldsymbol{k}}\tilde{\boldsymbol{y}}_{\boldsymbol{k}}$
	$\hat{\boldsymbol{P}}_{\boldsymbol{k}} = (\boldsymbol{I} - \boldsymbol{K}_{\boldsymbol{k}}\boldsymbol{H}_{\boldsymbol{k}})\overline{\boldsymbol{P}}_{\boldsymbol{k}}$

Figure 3*: The standard Kalman filter equations with the addition of the normalized innovation,* $\hat{y}_k$.

Two Monte Carlo simulations were run, varying the value of the modeled measurement error correlation parameter, $\rho \in \{0.9, 0.0\}$. The simulation parameters are tabulated in Table 1. In each simulation, Sequence and Sphericity monitors are initialized at the beginning of the simulation and permitted to accumulate innovations over the entire simulation horizon, engaging a hypothesis test after every measurement.

Figure 4 depicts the evaluated innovation test statistics for the two Monte Carlo scenarios over the simulation duration. The plot on the left depicts the Sequence and Sphericity test statistics and associated expected bounds under H_0 for a true-negative scenario; the plot on the right depicts the Sequence and Sphericity test statistics and associated bounds under H_0 for a true-positive scenario. The test statistic bounds are evaluated as bounds such that the respective cumulative chi-squared probability distributions evaluated at the bounds contain $1 -$ Prescribed False Postive Ratio of the probability.

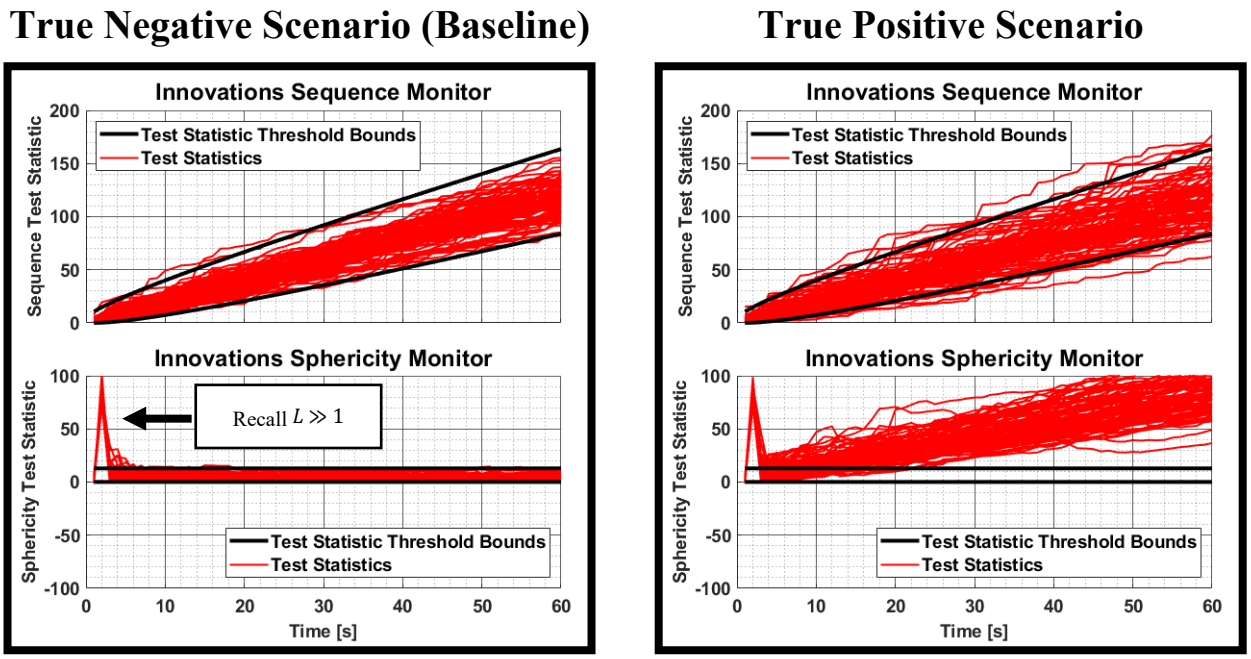

Figure 4: *Innovation test statistics for the two Monte Carlo scenarios. (Left) The test statistics and associated bounds under H_0 for the baseline scenario ($\rho^* = 0.9, \rho = 0.9$). (Right) The test statistics and associated bounds under H_0 for the true-positive scenario ($\rho^* = 0.9, \rho = 0.0$).*

The first simulation establishes a baseline: the navigation engineer correctly models the measurements as correlated ($\rho = 0.9$) when the measurement errors are in fact truly correlated ($\rho^* = 0.9$). This scenario establishes that both monitors operate within their prescribed false-positive ratio limits and establishes their true-negative asymptotic behaviors. Excepting an approximation-driven error in the Sphericity test statistics at the beginning of the simulation, nearly all the Monte Carlo test statistics (red) are contained within the expected test statistic bounds (black) – that is the monitors do not raise many false-positive flags. The brief foray of the Sphericity monitor test statistics at the beginning of the simulation is due to the requirement that the number of accumulated innovations, L, must be large to satisfy the chi-squared distribution approximation. By 10 seconds into the

# Monte Carlo	100	
Simulation Duration	60 s	
Initial Receiver Position	$[0, 0]^T$ m	
SV Positions	$[-30, 30]^T, [30, 30]^T$ m	
$\hat{\boldsymbol{x}}_{\boldsymbol{0}}$	$[0, 0]^T$ m	
$\hat{\boldsymbol{P}}_{\boldsymbol{0}}$	diag([10, 10]) m^2	
Δt	0.1 s	
$\boldsymbol{Q}$	diag([0.1, 0.1]) $(\text{m/s})^2/\text{Hz}$	
Measurement Frequency	1 Hz	
σ	3 m	
ρ^*	0.9	
ρ	0.9	0.0
False-Positive Ratio	0.01	

simulation (10 innovation samples), the chi-squared approximation becomes valid, and the test statistics become well-bounded by the expected bounds.

The second simulation mirrors that of the described scenario: the navigation engineer incorrectly models the measurements as uncorrelated ($\rho = 0$) when the measurement errors are in fact truly correlated ($\rho^* = 0.9$). This scenario establishes the behavior of the monitors under a true-positive scenario and evidences how the Sphericity monitor is effective at detecting correlation modeling errors while the Sequence monitor is not. Only a minor fraction – ~ 6% at any given time – of the test statistics for the Sequence monitor exceed the expected bounds under H_0 (a true positive), indicating that it is ineffective at detecting correlation modeling errors. In comparison, by 26 seconds into the simulation, all the Sphericity monitor test statistics had exceeded the expected bounds under H_0 and had raised a flag – it is extremely effective at detecting correlation modeling errors.

5. CONCLUSIONS & FUTURE WORK

Innovations Snapshot and Sequence monitors are demonstrated to be ineffective at detecting intra-epoch measurement correlation modeling errors. A new monitor, the innovations Sphericity monitor, is introduced which is specifically sensitive to such modeling errors. The effectiveness of this monitor is evidenced with a simple Gaussian sampling simulation and with a simple GNSS-like 2D Monte Carlo simulation. Future work may focus on (1) the sensitivity of the Sphericity monitor to other confounding modeling errors and non-linearities; (2) the Sphericity monitor's effectiveness when the dimension of the measurement is large ($M \gg 1$); (3) a comparison of the Sphericity monitor against other covariance-scrutinizing monitors; and (4) the application of the Sphericity monitor to specific navigation and GNSS problems.

6. ACKNOWLEDGEMENTS

This article has been authored by an employee of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employee owns all right, title and interest in and to the article and is solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan https://www.energy.gov/downloads/doe-public-access-plan.

This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

Approved for public release. SAND# XXX-XXXX

REFERENCES

Anderson, T. W. (1958). *An Introduction to Multivariate Statistical Analysis.* New York: John Wiley & Sons.

Giri, N. (1977). *Multivariate Statistical Inference.* New York: Academic Press Inc.

Groves, P. D. (2013). *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems, Second Edition.*

Hajiyev, C. (2008). Testing the Covariance Matrix of the Innovation Sequence in Application to Aircraft Sensor Fault Detection. *International federation of Automatic Control* . Seoul, Korea.

Hajiyev, C., & Caliskan, F. (1998). Fault Detection in Flight Control Systems via Innovation Sequence of Kalman Filter. *UKACC International Conference on Control.* Swansea.

Liu, Y., Li, S., Fu, Q., Liu, Z., & Zhou, Q. (2019). Analysis of Kalman Filter Innovation-Based GNSS Spoofing Detection Method for INS/GNSS Integrated Navigation System. *IEEE Sensors*, 5167-5178.

Mehra, R. K., & Peschon, J. (1971). An Innovations Approach to Fault Detection and Diagnosis in Dynamic Systems. *Automatica*, 637-640.

Muirhead, R. J. (2009). *Aspects of multivariate statistical theory.* John Wiley & Sons.

Quartararo, J. D., & Langel, S. E. (2021). Detecting Slowly Accumulating Faults Using a Bank of Cumulative Innovations Monitors in a Kalman Filter. *NAVIGATION*.

Tanil, C., Khanafseh, S., Joerger, M., & Pervan, B. (2016). Kalman filter-based INS monitor to detect GNSS spoofers cabable of tracking aircraft position. *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, (pp. 1027-1034).

Tanil, C., Khanafseh, S., Joerger, M., & Pervan, B. (2018). Sequential Integrity Monitoring for Kalman Filter Innovations-based Detectors. *ION GNSS+*. Miami.

APPENDIX A: INDEPENDENCE OF KALMAN FILTER INNOVATIONS

Let $\boldsymbol{x}_k^* \in \mathbb{R}^N$, $\bar{\boldsymbol{x}}_\boldsymbol{k} \in \mathbb{R}^N$, and $\bar{\boldsymbol{P}}_\boldsymbol{k} \in \mathbb{R}^{N\times N}$ be respectively the true state, estimated state, and associated estimated error covariance of a Kalman filter at time t_k. The estimated state is modeled as a random multivariate Gaussian vector centered on the true state with covariance $\bar{\boldsymbol{P}}_\boldsymbol{k}$:

$$\bar{\boldsymbol{x}}_\boldsymbol{k} \sim \boldsymbol{\mathcal{N}}\left(\boldsymbol{x}_\boldsymbol{k}^*, \bar{\boldsymbol{P}}_\boldsymbol{k}\right) \tag{23}$$

Suppose the state dynamics were modeled with the following linear (or sufficiently linearized) stochastic system:

$$\boldsymbol{x}_\boldsymbol{k}^* = \boldsymbol{F}_{\boldsymbol{k-1}}\boldsymbol{x}_{\boldsymbol{k-1}}^* + \boldsymbol{w}_{\boldsymbol{k-1}} \tag{24}$$

where $\boldsymbol{w}_{\boldsymbol{k-1}} \in \mathbb{R}^N$ is random propagation noise drawn from a zero-mean multivariate Gaussian distribution with covariance $\boldsymbol{Q}_{\boldsymbol{k-1}} \in \mathbb{R}^{N\times N}$, and $\boldsymbol{F}_{\boldsymbol{k-1}} \in \mathbb{R}^{N\times N}$ is the discrete-time state transition matrix.

Suppose the state were observed by a measurement $\boldsymbol{z}_\boldsymbol{k} \in \mathbb{R}^M$ modeled with the following linear (or sufficiently linearized) relationship:

$$\boldsymbol{z}_\boldsymbol{k} = \boldsymbol{H}_\boldsymbol{k}\boldsymbol{x}_\boldsymbol{k}^* + \boldsymbol{v}_\boldsymbol{k} \tag{25}$$

where $\boldsymbol{v}_\boldsymbol{k} \in \mathbb{R}^M$ is random measurement noise drawn from a zero-mean multivariate Gaussian distribution with covariance $\boldsymbol{R}_\boldsymbol{k} \in \mathbb{R}^{M\times M}$, and $\boldsymbol{H}_\boldsymbol{k} \in \mathbb{R}^{M\times N}$ is the linear observation map.

It is further assumed that the propagation and measurement noise samples are independent:

$$\boldsymbol{E}\left[\boldsymbol{w}_\boldsymbol{j}\boldsymbol{v}_\boldsymbol{k}^\boldsymbol{T}\right] = \boldsymbol{0}\ \forall\ j,k \tag{26}$$

The measurement innovation $\tilde{\boldsymbol{y}}_\boldsymbol{k} \in \mathbb{R}^M$ is defined:

$$\tilde{\boldsymbol{y}}_\boldsymbol{k} = \tilde{\boldsymbol{z}}_\boldsymbol{k} - \boldsymbol{H}_\boldsymbol{k}\bar{\boldsymbol{x}}_\boldsymbol{k} \tag{27}$$

where $\tilde{\boldsymbol{z}}_\boldsymbol{k} \in \mathbb{R}^M$ is the realized measurement at t_k.

Under the model assumptions in Equations (23) - (26), the measurement innovation should be zero mean:

$$\begin{aligned}\boldsymbol{E}[\tilde{\boldsymbol{y}}_\boldsymbol{k}] &= \mathrm{E}[\tilde{\boldsymbol{z}}_\boldsymbol{k} - \boldsymbol{H}_\boldsymbol{k}\bar{\boldsymbol{x}}_\boldsymbol{k}] \\ &= \mathrm{E}\left[\boldsymbol{H}_\boldsymbol{k}\boldsymbol{x}_\boldsymbol{k}^* + \boldsymbol{v}_\boldsymbol{k} - \boldsymbol{H}_\boldsymbol{k}\bar{\boldsymbol{x}}_\boldsymbol{k}\right] \\ &= \mathrm{E}\left[\boldsymbol{H}_\boldsymbol{k}\left(\boldsymbol{x}_\boldsymbol{k}^* - \bar{\boldsymbol{x}}_\boldsymbol{k}\right) + \boldsymbol{v}_\boldsymbol{k}\right] \\ &= \boldsymbol{H}_\boldsymbol{k}\,\mathrm{E}\left[\boldsymbol{x}_\boldsymbol{k}^* - \bar{\boldsymbol{x}}_\boldsymbol{k}\right] + \mathrm{E}[\boldsymbol{v}_\boldsymbol{k}] \\ &= \boldsymbol{0}\end{aligned} \tag{28}$$

Under the same assumptions, the variance of a measurement innovation should also be:

$$\begin{aligned}\mathrm{Var}[\tilde{\boldsymbol{y}}_\boldsymbol{k}] &= \mathrm{Var}[\tilde{\boldsymbol{z}}_\boldsymbol{k} - \boldsymbol{H}_\boldsymbol{k}\bar{\boldsymbol{x}}_\boldsymbol{k}] \\ &= \mathrm{Var}\left[\boldsymbol{H}_\boldsymbol{k}\boldsymbol{x}_\boldsymbol{k}^* + \boldsymbol{v}_\boldsymbol{k} - \boldsymbol{H}_\boldsymbol{k}\bar{\boldsymbol{x}}_\boldsymbol{k}\right] \\ &= \mathrm{Var}\left[\boldsymbol{H}_\boldsymbol{k}\left(\boldsymbol{x}_\boldsymbol{k}^* - \bar{\boldsymbol{x}}_\boldsymbol{k}\right) + \boldsymbol{v}_\boldsymbol{k}\right] \\ &= \mathrm{Var}\left[\boldsymbol{H}_\boldsymbol{k}\left(\boldsymbol{x}_\boldsymbol{k}^* - \bar{\boldsymbol{x}}_\boldsymbol{k}\right)\right] + \mathrm{Var}[\boldsymbol{v}_\boldsymbol{k}] \\ &= \boldsymbol{H}_\boldsymbol{k}\mathrm{Var}\left[\boldsymbol{x}_\boldsymbol{k}^* - \bar{\boldsymbol{x}}_\boldsymbol{k}\right]\boldsymbol{H}_\boldsymbol{k}^\mathrm{T} + \mathrm{Var}[\boldsymbol{v}_\boldsymbol{k}] \\ &= \underbrace{\boldsymbol{H}_\boldsymbol{k}\bar{\boldsymbol{P}}_\boldsymbol{k}\boldsymbol{H}_\boldsymbol{k}^\boldsymbol{T} + \boldsymbol{R}_\boldsymbol{k}}_{\boldsymbol{S}_\boldsymbol{k}}\end{aligned} \tag{29}$$

where $\boldsymbol{S}_\boldsymbol{k} \in \mathbb{R}^{M\times M}$ is the innovation covariance. The variance operator is distributed under the assumption that the measurement noise $\boldsymbol{v}_\boldsymbol{k}$ and the state estimate error $\boldsymbol{x}_\boldsymbol{k}^* - \bar{\boldsymbol{x}}_\boldsymbol{k}$ are independent.

To conclude, under the model assumptions in Equations (23) - (26), the innovations vector for a single measurement epoch $\tilde{\boldsymbol{y}}_{\boldsymbol{k}}$ should be distributed as:

$$\tilde{\boldsymbol{y}}_{\boldsymbol{k}} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{S}_{\boldsymbol{k}}) \tag{30}$$

If $\boldsymbol{S}_{\boldsymbol{k}}$ is decomposed such that $\boldsymbol{S}_{\boldsymbol{k}} = \tilde{\boldsymbol{S}}_{\boldsymbol{k}}\tilde{\boldsymbol{S}}_{\boldsymbol{k}}^{\boldsymbol{T}}$ (such as with a Cholesky decomposition), then it follows that the normalized innovation $\hat{\boldsymbol{y}}_{\boldsymbol{k}} \triangleq \tilde{\boldsymbol{S}}_{\boldsymbol{k}}^{-\boldsymbol{1}}\tilde{\boldsymbol{y}}_{\boldsymbol{k}}$ is distributed as a unit multivariate Gaussian distribution:

$$\hat{\boldsymbol{y}}_{\boldsymbol{k}} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}) \tag{31}$$

This is sufficient to prove that the subvectors of a normalized innovation are independent across a single measurement epoch but does not prove independence across several measurement epochs – this will be proven next.

Let $\boldsymbol{z}_{k+1} \in \mathbb{R}^M$ be a second observation produced by the same measurement system at some time in the future, t_{k+1}. Let $\boldsymbol{z}_{\boldsymbol{k+1}}$ be modeled similarly to $\boldsymbol{z}_{\boldsymbol{k}}$:

$$\boldsymbol{z}_{\boldsymbol{k+1}} = \boldsymbol{H}_{\boldsymbol{k+1}}\boldsymbol{x}_{\boldsymbol{k+1}}^{*} + \boldsymbol{v}_{\boldsymbol{k+1}} \tag{32}$$

where $\boldsymbol{v}_{\boldsymbol{k+1}} \in \mathbb{R}^M$ is random measurement noise drawn from a zero-mean multivariate Gaussian distribution with covariance $\boldsymbol{R}_{\boldsymbol{k+1}} \in \mathbb{R}^{M\times M}$, and $\boldsymbol{H}_{\boldsymbol{k+1}} \in \mathbb{R}^{M\times N}$ is the linear observation map. It is further assumed that this measurement noise is independent of any other process or measurement noise.

The future innovation vector $\tilde{\boldsymbol{y}}_{\boldsymbol{k+1}} \in \mathbb{R}^M$ is defined similarly to $\tilde{\boldsymbol{y}}_{\boldsymbol{k}}$:

$$\tilde{\boldsymbol{y}}_{\boldsymbol{k+1}} = \tilde{\boldsymbol{z}}_{\boldsymbol{k+1}} - \boldsymbol{H}_{\boldsymbol{k+1}}\overline{\boldsymbol{x}}_{\boldsymbol{k+1}} \tag{33}$$

where $\tilde{\boldsymbol{z}}_{\boldsymbol{k+1}} \in \mathbb{R}^M$ is the realized measurement at t_{k+1}.

The future prior state estimate $\overline{\boldsymbol{x}}_{\boldsymbol{k+1}}$ and future prior error covariance estimate $\overline{\boldsymbol{P}}_{\boldsymbol{k+1}} \in \mathbb{R}^{N\times N}$ are propagated from the current posterior state estimate $\hat{\boldsymbol{x}}_{\boldsymbol{k}} \in \mathbb{R}^N$ and corresponding error covariance $\hat{\boldsymbol{P}}_{\boldsymbol{k}} \in \mathbb{R}^{N\times N}$:

$$\begin{aligned} \overline{\boldsymbol{x}}_{\boldsymbol{k+1}} &= \boldsymbol{F}_{\boldsymbol{k}}\hat{\boldsymbol{x}}_{\boldsymbol{k}} \\ \overline{\boldsymbol{P}}_{\boldsymbol{k+1}} &= \boldsymbol{F}_{\boldsymbol{k}}\hat{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{F}_{\boldsymbol{k}}^{\boldsymbol{T}} + \boldsymbol{Q}_{\boldsymbol{k}} \end{aligned} \tag{34}$$

where $\hat{\boldsymbol{x}}_{\boldsymbol{k}}$ and $\hat{\boldsymbol{P}}_{\boldsymbol{k}}$ have been estimated from $\overline{\boldsymbol{x}}_{\boldsymbol{k}}$, $\overline{\boldsymbol{P}}_{\boldsymbol{k}}$, and $\tilde{\boldsymbol{z}}_{\boldsymbol{k}}$ using the Kalman update equations:

$$\begin{aligned} \tilde{\boldsymbol{y}}_{\boldsymbol{k}} &= \tilde{\boldsymbol{z}}_{\boldsymbol{k}} - \boldsymbol{H}_{\boldsymbol{k}}\overline{\boldsymbol{x}}_{\boldsymbol{k}} \\ \boldsymbol{S}_{\boldsymbol{k}} &= \boldsymbol{H}_{\boldsymbol{k}}\overline{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{H}_{\boldsymbol{k}}^{\boldsymbol{T}} + \boldsymbol{R}_{\boldsymbol{k}} \\ \boldsymbol{K}_{\boldsymbol{k}} &= \overline{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{H}_{\boldsymbol{k}}^{\boldsymbol{T}}\boldsymbol{S}_{\boldsymbol{k}}^{-\boldsymbol{1}} \\ \hat{\boldsymbol{x}}_{\boldsymbol{k}} &= \overline{\boldsymbol{x}}_{\boldsymbol{k}} + \boldsymbol{K}_{\boldsymbol{k}}\tilde{\boldsymbol{y}}_{\boldsymbol{k}} \\ \hat{\boldsymbol{P}}_{\boldsymbol{k}} &= (\boldsymbol{I} - \boldsymbol{K}_{\boldsymbol{k}}\boldsymbol{H}_{\boldsymbol{k}})\overline{\boldsymbol{P}}_{\boldsymbol{k}} \end{aligned} \tag{35}$$

Before evaluating the covariance of $\hat{\boldsymbol{y}}_{\boldsymbol{k}}$ and $\hat{\boldsymbol{y}}_{\boldsymbol{k+1}}$, the future measurement innovation $\tilde{\boldsymbol{y}}_{\boldsymbol{k+1}}$ is reformulated in terms of $\boldsymbol{x}_{\boldsymbol{k}}^{*}$, $\overline{\boldsymbol{x}}_{\boldsymbol{k}}$, and $\overline{\boldsymbol{P}}_{\boldsymbol{k}}$:

$$\begin{aligned} \tilde{\boldsymbol{y}}_{\boldsymbol{k+1}} &= \tilde{\boldsymbol{z}}_{\boldsymbol{k+1}} - \boldsymbol{H}_{\boldsymbol{k+1}}\overline{\boldsymbol{x}}_{\boldsymbol{k+1}} \\ &= \boldsymbol{H}_{\boldsymbol{k+1}}\boldsymbol{x}_{\boldsymbol{k+1}}^{*} + \boldsymbol{v}_{\boldsymbol{k+1}} - \boldsymbol{H}_{\boldsymbol{k+1}}\overline{\boldsymbol{x}}_{\boldsymbol{k+1}} \\ &= \boldsymbol{H}_{\boldsymbol{k+1}}\left(\boldsymbol{F}_{\boldsymbol{k}}\boldsymbol{x}_{\boldsymbol{k}}^{*} + \boldsymbol{w}_{\boldsymbol{k}}\right) + \boldsymbol{v}_{\boldsymbol{k+1}} - \boldsymbol{H}_{\boldsymbol{k+1}}\boldsymbol{F}_{\boldsymbol{k}}\hat{\boldsymbol{x}}_{\boldsymbol{k}} \\ &= \boldsymbol{H}_{\boldsymbol{k+1}}\left(\boldsymbol{F}_{\boldsymbol{k}}\boldsymbol{x}_{\boldsymbol{k}}^{*} + \boldsymbol{w}_{\boldsymbol{k}}\right) + \boldsymbol{v}_{\boldsymbol{k+1}} - \boldsymbol{H}_{\boldsymbol{k+1}}\boldsymbol{F}_{\boldsymbol{k}}(\overline{\boldsymbol{x}}_{\boldsymbol{k}} + \boldsymbol{K}_{\boldsymbol{k}}\tilde{\boldsymbol{y}}_{\boldsymbol{k}}) \\ &= \boldsymbol{H}_{\boldsymbol{k+1}}\left(\boldsymbol{F}_{\boldsymbol{k}}\boldsymbol{x}_{\boldsymbol{k}}^{*} + \boldsymbol{w}_{\boldsymbol{k}}\right) + \boldsymbol{v}_{\boldsymbol{k+1}} - \boldsymbol{H}_{\boldsymbol{k+1}}\boldsymbol{F}_{\boldsymbol{k}}(\overline{\boldsymbol{x}}_{\boldsymbol{k}} + \boldsymbol{K}_{\boldsymbol{k}}(\tilde{\boldsymbol{z}}_{\boldsymbol{k}} - \boldsymbol{H}_{\boldsymbol{k}}\overline{\boldsymbol{x}}_{\boldsymbol{k}})) \\ &= \boldsymbol{H}_{\boldsymbol{k+1}}\left(\boldsymbol{F}_{\boldsymbol{k}}\boldsymbol{x}_{\boldsymbol{k}}^{*} + \boldsymbol{w}_{\boldsymbol{k}}\right) + \boldsymbol{v}_{\boldsymbol{k+1}} - \boldsymbol{H}_{\boldsymbol{k+1}}\boldsymbol{F}_{\boldsymbol{k}}\left(\overline{\boldsymbol{x}}_{\boldsymbol{k}} + \boldsymbol{K}_{\boldsymbol{k}}\left(\boldsymbol{H}_{\boldsymbol{k}}\boldsymbol{x}_{\boldsymbol{k}}^{*} + \boldsymbol{v}_{\boldsymbol{k}} - \boldsymbol{H}_{\boldsymbol{k}}\overline{\boldsymbol{x}}_{\boldsymbol{k}}\right)\right) \end{aligned} \tag{36}$$

Distributing and collecting like terms yields:

$$\tilde{\boldsymbol{y}}_{\boldsymbol{k+1}} = \boldsymbol{H_{k+1}F_k}(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}}) - \boldsymbol{H_{k+1}F_kK_kH_k}(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}}) + \boldsymbol{H_{k+1}w_k} + \boldsymbol{v_{k+1}} - \boldsymbol{H_{k+1}F_kK_kv_k} \quad (37)$$

Now, the covariance of $\hat{\boldsymbol{y}}_{\boldsymbol{k}}$ and $\hat{\boldsymbol{y}}_{\boldsymbol{k+1}}$ may be readily evaluated:

$$\begin{aligned}
\boldsymbol{Cov}[\hat{\boldsymbol{y}}_{\boldsymbol{k}}, \hat{\boldsymbol{y}}_{\boldsymbol{k+1}}] &= \boldsymbol{E}\left[\hat{\boldsymbol{y}}_{\boldsymbol{k}}\hat{\boldsymbol{y}}_{\boldsymbol{k+1}}^{\boldsymbol{T}}\right] \\
&= \boldsymbol{E}\Big[\left(\boldsymbol{H_k}(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}}) + \boldsymbol{v_k}\right) \cdot \left(\boldsymbol{H_{k+1}F_k}(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}}) - \boldsymbol{H_{k+1}F_kK_kH_k}(\boldsymbol{x_k^*} - \right. \\
&= \boldsymbol{E}\Big[\boldsymbol{H_k}(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}})(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}})^{\boldsymbol{T}}\boldsymbol{F_k^TH_{k+1}^T} - \boldsymbol{H_k}(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}})(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}})^{\boldsymbol{T}}\boldsymbol{H_k^TK_k^TF_k^TH_{k+1}^T} + \boldsymbol{H_k}(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}})\boldsymbol{w_k^TH_{k+1}^T} + \boldsymbol{H_k} \\
&\quad (\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}})\boldsymbol{v_{k+1}^T} - \boldsymbol{H_k}(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}})\boldsymbol{v_k^TK_k^TF_k^TH_{k+1}^T} + \boldsymbol{v_k}(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}})^{\boldsymbol{T}}\boldsymbol{F_k^TH_{k+1}^T} - \boldsymbol{v_k}(\boldsymbol{x_k^*} - \bar{\boldsymbol{x}}_{\boldsymbol{k}})^{\boldsymbol{T}}\boldsymbol{H_k^TK_k^TF_k^TH_{k+1}^T} \\
&\quad + \boldsymbol{v_kw_k^TH_{k+1}^T} + \boldsymbol{v_kv_{k+1}^T} - \boldsymbol{v_kv_k^TK_k^TF_k^TH_{k+1}^T}\Big] \\
&= \boldsymbol{H_k}\bar{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{F_k^TH_{k+1}^T} - \boldsymbol{H_k}\bar{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{H_k^TK_k^TF_k^TH_{k+1}^T} - \boldsymbol{R_kK_k^TF_k^TH_{k+1}^T} \\
&= \boldsymbol{H_k}\bar{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{F_k^TH_{k+1}^T} - \left(\boldsymbol{H_k}\bar{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{H_k^T} + \boldsymbol{R_k}\right)\boldsymbol{K_k^TF_k^TH_{k+1}^T} \\
&= \boldsymbol{H_k}\bar{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{F_k^TH_{k+1}^T} - \boldsymbol{S_kK_k^TF_k^TH_{k+1}^T} \\
&= \boldsymbol{H_k}\bar{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{F_k^TH_{k+1}^T} - \boldsymbol{H_k}\bar{\boldsymbol{P}}_{\boldsymbol{k}}\boldsymbol{F_k^TH_{k+1}^T} \\
&= \boldsymbol{0}
\end{aligned}$$

The covariance is zero and implies statistical independence between the innovations. This result shows that under several linear, Gaussian, and independence assumptions, normalized measurement innovations should be independent across a single measurement epoch and between measurement epochs.

APPENDIX B: DERIVATION OF SPHERICITY TEST STATISTIC

The following appendix derives the sphericity test statistic. It is based on the proofs established in Anderson (1958), Giri (1977), and Muirhead (2009) which are more rigorous, include minor adjustments to account for sample covariance biases, and derive the asymptotic approximation error distributions.

Suppose a Kalman filter produces normalized measurement innovations over N measurement epochs. Let $\hat{\boldsymbol{y}}_{\boldsymbol{k}} \in \mathbb{R}^p$ denote the kth measurement innovation. Suppose the normalized innovations were modeled as independently sampled from a multivariate Gaussian distribution with mean $\boldsymbol{\mu} \in \mathbb{R}^p$ and covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{p\times p}$. Consider testing the null hypothesis H_0:$\boldsymbol{\Sigma} = \boldsymbol{I}$ against the alternate hypothesis H_A:$\boldsymbol{\Sigma} \neq \boldsymbol{I}$. Muirhead describes this hypothesis as the 'hypothesis of sphericity,' as the contours of equal probability density under the null hypothesis form hyperspheres – hence the same 'sphericity test'.

The likelihood of a single multivariate Gaussian observation $\hat{\boldsymbol{y}} \in \mathbb{R}^p$ with mean $\boldsymbol{\mu} \in \mathbb{R}^p$ and covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{p\times p}$ is:

$$L(\boldsymbol{\mu},\boldsymbol{\Sigma}) = |2\pi\boldsymbol{\Sigma}|^{-1/2}\exp\left(-\frac{1}{2}(\hat{\boldsymbol{y}} - \boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\hat{\boldsymbol{y}} - \boldsymbol{\mu})\right) \quad (38)$$

where $|\cdot|$ is the matrix determinant.

It follows that the likelihood of N independent samples from the same distribution is:

$$L(\boldsymbol{\mu},\boldsymbol{\Sigma}) = |2\pi\boldsymbol{\Sigma}|^{-N/2}\prod_{i=1}^{N}\exp\left(-\frac{1}{2}(\hat{\boldsymbol{y}}_i - \boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\hat{\boldsymbol{y}}_i - \boldsymbol{\mu})\right) \quad (39)$$

The product of exponentials is elevated into the exponential as a sum:

$$L(\boldsymbol{\mu},\boldsymbol{\Sigma}) = |2\pi\boldsymbol{\Sigma}|^{-N/2}\exp\left(-\frac{1}{2}\sum_{i=1}^{N}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})\right) \tag{40}$$

Employing a matrix trace trick to isolate the mean and covariance components:

$$\begin{aligned}
L(\boldsymbol{\mu},\boldsymbol{\Sigma}) &= |2\pi\boldsymbol{\Sigma}|^{-N/2}\exp\left(-\frac{1}{2}\sum_{i=1}^{N}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})\right)\\
&= |2\pi\boldsymbol{\Sigma}|^{-N/2}\exp\left(-\frac{1}{2}\mathrm{tr}\left[\sum_{i=1}^{N}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})\right]\right)\\
&= |2\pi\boldsymbol{\Sigma}|^{-N/2}\exp\left(-\frac{1}{2}\sum_{i=1}^{N}\mathrm{tr}\left[(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})\right]\right)\\
&= |2\pi\boldsymbol{\Sigma}|^{-N/2}\exp\left(-\frac{1}{2}\sum_{i=1}^{N}\mathrm{tr}\left[\boldsymbol{\Sigma}^{-1}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})^T\right]\right)\\
&= |2\pi\boldsymbol{\Sigma}|^{-N/2}\exp\left(-\frac{1}{2}\mathrm{tr}\left[\sum_{i=1}^{N}\boldsymbol{\Sigma}^{-1}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})^T\right]\right)\\
&= |2\pi\boldsymbol{\Sigma}|^{-N/2}\exp\left(-\frac{1}{2}\mathrm{tr}\left[\boldsymbol{\Sigma}^{-1}\sum_{i=1}^{N}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})^T\right]\right)
\end{aligned} \tag{41}$$

Next, a likelihood ratio $\Lambda \in \mathbb{R}$ is formed:

$$\Lambda = \frac{L_{H_0}(\boldsymbol{\mu},\boldsymbol{\Sigma}=\mathbf{I})}{L_{H_A}(\boldsymbol{\mu},\boldsymbol{\Sigma}\neq\mathbf{I})} = \frac{|2\pi\boldsymbol{I}|^{-N/2}\exp\left(-\frac{1}{2}\mathrm{tr}\left[\boldsymbol{I}^{-1}\sum_{i=1}^{N}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})^T\right]\right)}{|2\pi\boldsymbol{\Sigma}|^{-N/2}\exp\left(-\frac{1}{2}\mathrm{tr}\left[\boldsymbol{\Sigma}^{-1}\sum_{i=1}^{N}(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})(\hat{\boldsymbol{y}}_i-\boldsymbol{\mu})^T\right]\right)} \tag{42}$$

The parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are unknown, so a generalized likelihood ratio must be employed. A generalized likelihood ratio replaces the unknown parameters with their respective maximum-likelihood estimates under the respective hypotheses. It is straightforward to verify that the maximum-likelihood estimates of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are the sample mean and covariance – that is, given data with sample mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$**,** the most likely values of the distribution mean and covariance are the sample mean and covariance.

Let $\overline{\boldsymbol{y}} \in \mathbb{R}^p$ be the sample mean and $\boldsymbol{S} \in \mathbb{R}^{p\times p}$ be the sample covariance of $\hat{\boldsymbol{y}}_{1..N}$:

$$\overline{\boldsymbol{y}} = \frac{1}{N}\sum_{i=1}^{N}\hat{\boldsymbol{y}}_{\boldsymbol{i}} \tag{43}$$

$$\boldsymbol{S} = \frac{1}{N}\sum_{i=1}^{N}(\hat{\boldsymbol{y}}_{\boldsymbol{i}}-\overline{\boldsymbol{y}})(\hat{\boldsymbol{y}}_{\boldsymbol{i}}-\overline{\boldsymbol{y}})^T \tag{44}$$

Furthermore, let $\boldsymbol{B} \in \mathbb{R}^{p\times p}$ be the scaled sample covariance:

$$\boldsymbol{B} = N\boldsymbol{S} = \sum_{i=1}^{N}(\hat{\boldsymbol{y}}_{\boldsymbol{i}}-\overline{\boldsymbol{y}})(\hat{\boldsymbol{y}}_{\boldsymbol{i}}-\overline{\boldsymbol{y}})^T \tag{45}$$

Substituting $\boldsymbol{\mu} = \overline{\boldsymbol{y}}$ and $\boldsymbol{\Sigma} = \boldsymbol{S}$ into the likelihood ratio yields:

$$\Lambda = \frac{L_{H_0}(\boldsymbol{\mu}=\overline{\boldsymbol{y}},\boldsymbol{\Sigma}=\mathbf{I})}{L_{H_A}(\boldsymbol{\mu}=\overline{\boldsymbol{y}},\boldsymbol{\Sigma}=\mathbf{S})} = \frac{|2\pi\boldsymbol{I}|^{-N/2}\exp\left(-\frac{1}{2}\mathrm{tr}[\boldsymbol{I}^{-1}\boldsymbol{B}]\right)}{|2\pi\boldsymbol{S}|^{-N/2}\exp\left(-\frac{1}{2}\mathrm{tr}[\boldsymbol{S}^{-1}\boldsymbol{B}]\right)} \tag{46}$$

Simplifying and noting that $\boldsymbol{S} = \boldsymbol{B}/N$:

$$\Lambda = \left(\frac{e}{N}\right)^{Np/2} |\boldsymbol{B}|^{N/2} \exp\left(-\frac{1}{2}\mathrm{tr}[\boldsymbol{B}]\right) \tag{47}$$

Further simplifying:

$$\begin{aligned} \Lambda^* &= -2\ln(\Lambda) \\ &= -Np\big(1-\ln(N)\big) - N\ln\big(\det(\boldsymbol{B})\big) + \mathrm{tr}[\boldsymbol{B}] \end{aligned} \tag{48}$$

Finally, the generalized maximum-likelihood hypothesis test may be stated:

$$\underbrace{-Np\big(1-\ln(N)\big) - N\ln\big(\det(\boldsymbol{B})\big) + \mathrm{tr}[\boldsymbol{B}]}_{\Lambda^*} \underset{H_0}{\overset{H_A}{\gtrless}} \Lambda_0 \tag{49}$$

The hypothesis test will reject Λ^* whenever it is greater than a constant Λ_0 which is chosen such that the test's false-positive ratio is $\alpha \in \mathbb{R}$. To evaluate Λ_0, the distribution of Λ^* under H_0 must be derived. Here the derivation follows that of Giri (1977).

Under the null hypothesis that $\hat{\boldsymbol{y}} \sim \mathcal{N}(\overline{\boldsymbol{y}}, \boldsymbol{I})$, the scaled sample covariance matrix $\boldsymbol{B} = \sum_{i=1}^{N}(\hat{\boldsymbol{y}}_{\boldsymbol{i}} - \overline{\boldsymbol{y}})(\hat{\boldsymbol{y}}_{\boldsymbol{i}} - \overline{\boldsymbol{y}})^T$ has a Wishart distribution with scale parameter $\boldsymbol{I}$ and $N-1$ degrees of freedom. Anderson (1958) provides the characteristic function of Λ^*, $\phi(\Lambda^*)$, as the product:

$$\phi(\Lambda^*) = \prod_{j=1}^{p} \phi_j(t) \tag{50}$$

where

$$\phi_j(t) = (2e/N)^{-iNt}(1-2it)^{-(N-1-2iNt)/2} \frac{\Gamma\left(\frac{1}{2}(N-j) - iNt\right)}{\Gamma\left(\frac{1}{2}(N-j)\right)} \tag{51}$$

Substituting Stirling's approximation for the Gamma function

$$\Gamma(z) = \left(\frac{2\pi}{z}\right)^{\frac{1}{2}} \left(\frac{z}{e}\right)^{z} \tag{52}$$

into Equation (51) approximates the characteristic function as:

$$\phi_j(t) \simeq (2e/N)^{-iNt}(1-2it)^{-(N-1-2iNt)/2} \frac{\left(\frac{2\pi}{\frac{1}{2}(N-j) - iNt}\right)^{\frac{1}{2}} \left(\frac{\frac{1}{2}(N-j) - iNt}{e}\right)^{\frac{1}{2}(N-j)-iNt}}{\left(\frac{2\pi}{\frac{1}{2}(N-j)}\right)^{\frac{1}{2}} \left(\frac{\frac{1}{2}(N-j)}{e}\right)^{\frac{1}{2}(N-j)}} \tag{53}$$

Simplifying:

$$\begin{aligned}
\phi_j(t) &\simeq (1-2it)^{-\frac{1}{2}(N-1-2iNt)}\left(\frac{N-j}{N-j-2iNt}\right)^{\frac{1}{2}}\left(\frac{\frac{1}{2}(N-j)-iNt}{\frac{1}{2}(N-j)}\right)^{\frac{1}{2}(N-j)}\left(\frac{N-j-2iNt}{N}\right)^{-iNt} \\
&\simeq (1-2it)^{-\frac{1}{2}(N-1)}\cdot\left(\frac{N-j-2iNt}{N-j}\right)^{\frac{1}{2}(N-j-1)}\left(\frac{N-j-2iNt}{N(1-2it)}\right)^{-iNt} \\
&\simeq \frac{(1-2it)^{-j/2}}{(1-2it)^{-j/2}}(1-2it)^{-\frac{1}{2}(N-1)}\cdot\left(\frac{N-j-2iNt}{N-j}\right)^{\frac{1}{2}(N-j-1)}\left(\frac{N-j-2iNt}{N(1-2it)}\right)^{-iNt} \\
&\simeq (1-2it)^{-j/2}\left(\frac{N-j-2iNt}{(N-j)(1-2it)}\right)^{\frac{1}{2}(N-j-1)}\left(\frac{N-j-2iNt}{N(1-2it)}\right)^{-iNt} \\
&\simeq (1-2it)^{-j/2}\cdot\left(\frac{N(1-2it)-j}{N(1-2it)-j(1-2it)}\right)^{\frac{1}{2}(N-j-1)}\cdot\left(1-\frac{j}{N(1-2it)}\right)^{-iNt}
\end{aligned} \tag{54}$$

Then, as $N\to\infty$, the second and third terms reduce to unity and leave:

$$\phi_j(t) \simeq (1-2it)^{-j/2} \tag{55}$$

Thus, the asymptotic characteristic function for Λ^* is:

$$\phi(\Lambda^*) = \prod_{j=1}^{p}(1-2it)^{-j/2} = (1-2it)^{-p(p+1)/2} \tag{56}$$

which is the characteristic function of a chi-squared distribution with $p(p+1)/2$ degrees of freedom.

To summarize, under the null hypothesis, as $N\to\infty$, $\Lambda^* = -Np\left(1-\ln(N)\right) - N\ln\left(\det(\boldsymbol{B})\right) + \mathrm{tr}[\boldsymbol{B}]$ should be distributed as a chi-squared random variable with $p(p+1)/2$ degrees of freedom. The characterization of this distribution under H_0 is sufficient to determine Λ_0 provided an allowable, prescribed false-positive ratio of α.