Optimal Sequential Spoof Detection based on Direction of Arrival Measurements

Fabian Rothmaier, *Stanford University*

BIOGRAPHY

Fabian Rothmaier is a PhD candidate at the GPS Laboratory at Stanford University. He received his B. Engr. degree from the University of Applied Sciences Bremen, Germany in 2015 and his M. Sc. degree from Stanford University in 2017.

ABSTRACT

In this paper we present a sequential likelihood ratio test to detect a GNSS spoofing attack using information from multiple measurements. We develop an analytical solution for the test's detection threshold that can be computed online by the receiver and that is optimal according to the Neyman-Pearson lemma. We present its application to both independent and correlated Direction of Arrival measurements. To demonstrate the procedure's general applicability, we reproduce previously published simulations and apply it to data from flight tests and a live spoofing event.

I. INTRODUCTION

A lot of work has been done in recent years on GNSS interference, specifically spoofing detection and mitigation. Approaches within the GNSS receiver can roughly be categorized into spatial processing approaches that leverage the spatial diversity in the received signals, and signal quality based techniques that monitor characteristics of each individual received signal to detect anomalies. A general overview of the wide range of proposed techniques is given by several publications such as [1].

Spatial processing approaches usually require additional hardware compared to a traditional antenna-receiver setup to generate a spatial metric. This can be realized through e.g. a moving antenna as in [2], multiple separate antennas as in [3] or [4], antenna arrays as in [5] or [6] or a single element Dual Polarization Antenna (DPA) as in [7], [8]. The spatial metric, often referred to as the signal's Direction of Arrival (DoA), is then used for spoof detection. The underlying assumption is that an attacker would broadcast at least a subset of satellite signals from the same direction while the authentic signals are coming from a diverse set of directions. The majority of the paper deals with this type of approach and the underlying procedure and hypothesis will be explained in more detail in the following sections.

Radio Power Monitoring techniques commonly examine use of an incoming power monitor such as Automatic Gain Control (AGC) in the receiver front end [9] along with the monitoring of the received signal power relative to noise such as carrier to noise ration (C/N0).

Signal Quality Monitoring (SQM) techniques examine different metrics on the correlation function to determine the presence of a spoof signal mixed within the genuine signal. Examples of such techniques are given by [10], [11] and as shown e.g. by the group that published [12]. Other SQM techniques look for the presence of a second signal outside the normal correlation window to find steady state spoofing [13].

Every approach has its drawbacks, and an attack scenario can be conceived that overcomes each defense. Jamming of the receiver at the beginning of the attack will likely overcome many SQM based techniques. Spatial processing based approaches reach their limitation when only a few and closely spaced satellite signals are spoofed from a single source as illustrated in [5]. They have further yet to be proven in challenging urban environments, where the occlusion of large parts of the sky by buildings reduces the available satellite geometry and multipath distorts the signal's DoAs.

Some work has been done to combine different metrics for detection to overcome some of these limitations. For example, in [14] the authors argue that signal power and information about the correlation function complement each other nicely. They show promising results by tightly coupling the two metrics for a decision. In [6] monitoring of the pseudorange residuals are combined with DoA measurements. Scenarios where only a subset of satellites is spoofed and the receiver still receives some authentic signals are challenging for DoA based approaches, as fewer signals received from the same direction is less telling. This will however result in a disagreement between the pseudoranges to authentic and spoofed satellites, leading to elevated residuals. Detecting an attack due to increased pseudorange residuals potentially works well in a situation where DoA based approaches struggle, making this a prime example of different techniques complementing each other.

Not all metrics are this complementary however, and it is not straight forward to make a statistically justified decision about the presence of an attack based on multiple techniques. A receiver however has access to multiple metrics (pseudorange residuals, AGC, correlation function), and if one of the mentioned antenna modifications is employed spatial DoA information is available as well. The most educated decision about the presence of a spoofing attack or interference in general will leverage all information available within a certain time frame, making a successful attack as difficult and costly as possible.

In this paper we present a Likelihood Ratio Test (LRT) for a batch of measurements. It combines information from multiple metrics or from the same metric at multiple epochs. As a first application example, we combine DoA measurements from multiple sequential epochs for both the case of independent and correlated measurements. We demonstrate an increase in test power and number of detections under conditions challenging for DoA based approaches, increasing the capability of this type of spoofing defense.

Overall this paper makes three contributions towards mitigating the limitations of a single spoof detection technique.
First, we expand a LRT framework that we presented in [15] for batches of sequential measurements. We phrase hypothesis that result in a Normally distributed decision variable, leading to a straightforward analytic solution of the detection threshold that can be computed online by the receiver. The test guarantees a chosen false alert probability and is optimal according to the Neyman-Pearson lemma [16].
As a second contribution, we demonstrate the application to a batch of independent DoA measurements. We emphasize the low computational complexity and general applicability of the LRT by reproducing previously published simulations. We make our code publicly available for transparency and reproducibility.
Lastly, we demonstrate the application to a batch of correlated azimuth only DoA measurements. We characterize the correlation coefficient for an antenna mounted on an aircraft in flight and show favorable behavior towards the constraint on false alerts. We finally apply the sequential LRT to correlated sequential data from a government sponsored live spoofing event and demonstrate an increase in theoretical test power and actual number of detections.

The remaining paper is divided into three sections plus a summary and conclusions. In section II we set up the LRT for a batch of DoA measurements. In section III we consider the case of independent measurements. In section IV we consider correlated measurements. We analyze nominal data collected in flight and spoofed data collected during a government sponsored event.

II. THE SEQUENTIAL SPOOF DETECTION FRAMEWORK

A sophisticated spoofing attack likely needs to persist for some time until it results in hazardous misleading information. To avoid triggering simple detectors, like checks on the physical behavior of a vehicle, the attacker will only slowly alter the victim's navigation solution. The receiver therefore likely gathers multiple measurements of e.g. DoAs until any harm is caused by an attack. For a more informed decision about the presence of an attack, we propose to leverage information from these multiple measurements.

The approach to exploit the timely relation between multiple epochs is already used in [17] by simply counting consecutive threshold exceedances within the TTA, a concept similar to Innovation Sequence Monitoring [18]. The idea of a sequential filter is more specifically mentioned e.g. in [19], but to the author's best knowledge never exploited to track the probability of a spoofing attack.

To conceptually illustrate the difference between snapshot and sequential spoof detection, we depict graphical models of the snapshot based detection in Figure 1 and of sequential detection in the case of independent measurements in Figure 2. The state X is indicated in blue and the measurement or evidence Y in yellow for three epochs t_1, t_2 and t_3.

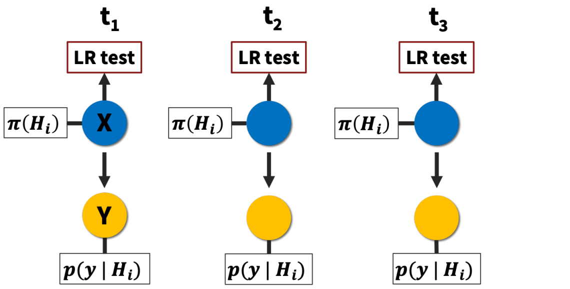

Figure 1: Graphical model of snapshot-based spoof detection for three epochs. The state X is depicted in blue, evidence Y in yellow. Prior probabilities $\pi(H_i)$ are not affecting the LRT but are depicted to match the standard representation of a Graphical model.

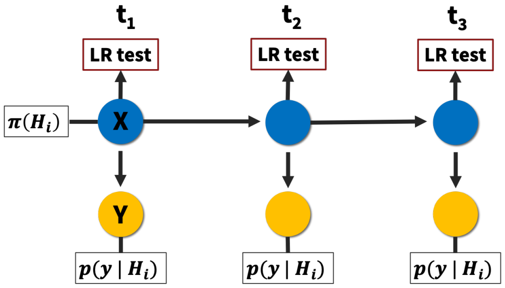

Figure 2: Graphical model of sequential spoof detection for three epochs if the measurements can be considered independent. The state X is depicted in blue, evidence Y in yellow. The prior probability at one epoch is the posterior probability of the preceding epoch.

In the snapshot approach, at each epoch the measurement and prior probabilities lead to a belief about the state. A decision is then made e.g. through a hypothesis or LRT. While prior probabilities $\pi(H_i)$ have no effect in the LRT following the Neyman-

Pearson paradigm used in this paper, we depict them here to match the standard representation of a Bayesian Network. In a sequential approach using independent measurements, the prior probabilities are based on the posterior belief of the previous epoch [20]. The LRT at time t_3 is based on evidence gathered during all three epochs t_1, t_2 and t_3.

In this section we derive the Likelihood Ratio Test (LRT) for a general batch of measurements, an analytic formulation of the detection threshold and the test power. In the following two sections we then consider the application to batches of sequential independent and correlated measurements.

1. Detection as a Hypothesis Test

Any decision about the presence of a spoofing attack is a decision under uncertainty. We therefore start by casting the decision problem as a statistical hypothesis test as we have in [8]. The null hypothesis, from here on denoted H_0, represents the nominal situation without spoofing. The alternate hypothesis H_1 represents a spoofed situation.

The prior probability of a spoofing attack is difficult to estimate and can vary greatly with time and place of the application. We therefore follow the Neyman-Pearson paradigm that is independent of prior probabilities as we will see in the following steps and it is shown in e.g. [21].
The hypothesis test is the solution to an optimization problem: We maximize detections while satisfying constraints on the false alert probability or statistical significance level P_{FAmax} and the Time To Alert (TTA) in Equation (1). An "alert" is equal to the rejection of H_0.

$$\begin{aligned} \max \quad & \textit{detections} \\ s.t. \quad & p(\textit{false alert}) \leq P_{FA_{max}} \\ & TTA \leq TTA_{max} \end{aligned} \tag{1}$$

In this paper we apply the hypothesis test to a batch of measurements. A hypothesis test based on a mix of genuine and spoofed signals is less powerful than a test based on evidence indicate the same state. To guarantee maximum detection power while satisfying the constraint on TTA_{max}, we only take into account measurements more recent than the maximum TTA for the specific application. In other words a decision at time t_k after a required processing time t_{proc} can only be based on measurement between epochs $t_1, \ldots, t_k$ such that the inequality $t_k - t_1 \leq TTA_{max} - t_{proc}$ is satisfied.

H_0 is rejected if the probability of H_1 given some measurements between epochs $t_1, \ldots, t_k$ collected in the vector $\boldsymbol{y_{t_1:t_k}}$ is above some threshold c, resulting in the formulation of the false alert constraint in Eq. (2). We developed the detailed derivation of the LRT in [8] and repeat the major steps here for convenience. Using Bayes Rule and indicating the prior probability of the ith hypothesis by $\pi(H_i)$, we modify the constraint equation and collect Random Variables (RVs) on one side. It is further a conservative choice to replace the joint event of H_0 with a conditioning on H_0, leading to the well-known formulation of a LRT under the Neyman-Pearson paradigm [21].

$$p\left(p\left(H_1|\boldsymbol{y_{t_1:t_k}}\right) > c, H_0\right) \leq P_{FA_{max}} \tag{2}$$

$$p\left(\frac{p\left(\boldsymbol{y_{t_1:t_k}}|H_0\right)}{p\left(\boldsymbol{y_{t_1:t_k}}|H_1\right)} < \left(\frac{1}{c} - 1\right)\frac{\pi(H_1)}{\pi(H_0)}|H_0\right) \leq P_{FA_{max}} \tag{3}$$

The only RV in Equation (3) is the likelihood ratio on the left-hand side of the inequality which we denote as $\Lambda\left(\mathrm{y}_{\boldsymbol{t_1:t_k}}\right)$. We collapse all constants on the right-hand side of the inequality into the detection threshold γ, marking the procedure's independence of the prior probabilities. As we will see later on, it is desirable to work with the natural log of both sides, leading to the formulation of the constraint given by Eq. (4).

$$p\left(\log\Lambda\left(\mathrm{y}_{\boldsymbol{t_1:t_k}}\right) < \log\gamma|H_0\right) \leq P_{FA_{max}} \quad \text{with} \quad \Lambda\left(\mathrm{y}_{\boldsymbol{t_1:t_k}}\right) = \frac{p\left(\boldsymbol{y_{t_1:t_k}}|H_0\right)}{p\left(\boldsymbol{y_{t_1:t_k}}|H_1\right)} \tag{4}$$

This represents the most powerful test according to the Neyman-Pearson lemma [16]. The null hypothesis is rejected if $\log\Lambda\left(\mathrm{y}_{\boldsymbol{t_1:t_k}}\right)$ is below the threshold $\log\gamma$. The largest γ satisfying Eq. (4) is the solution to the original optimization problem formulated in Equation (1).
$\log\gamma$ is found by solving the quantile function or inverse cumulative density function (cdf) of the random variable $\log\Lambda\left(\mathrm{y}_{\boldsymbol{t_1:t_k}}\right)$

with $\boldsymbol{y_{t_1:t_k}} \sim H_0$ for $P_{FA_{max}}$. This can be difficult to solve analytically and might have to be done with the help of Monte-Carlo analysis run offline as done in [3], [8] for snapshot based detection. This quickly becomes computationally intractable for the sequential setup. In the next subsection we phrase hypotheses for Direction of Arrival (DoA) based approaches such that $\log \Lambda \left(\mathrm{y}_{\boldsymbol{t_1:t_k}}\right)$ is a Normally distributed RV, leading to a general analytic solution for the detection threshold that can be calculated online by the receiver. We will then expand the formulation for batches of both independent and correlated sequential measurements.

2. Derivation of the Likelihood Ratio for Direction of Arrival Measurements

In the past subsection we have summarized the general LRT formulation for some measurements $\boldsymbol{y_{t_1:t_k}}$. We will now specify how to calculate $\log \Lambda \left(\mathrm{y}_{\boldsymbol{t_1:t_k}}\right)$ and the detection threshold $\log \gamma$ for spoof detection approaches working with the Direction of Arrival (DoA) of signals. In the following sections we will then consider specifically both the case of independent and correlated measurements.

In DoA based spoof detection, H_1 is generally specified as all or a subset of signals coming from the same direction as in references [5], [6], [22]. In [8] we have shown how an attack can be identified efficiently if only a subsets of the satellite signals is received from the same direction by using a iterative sequence of "all nominal" vs. "all spoofed" tests. We therefore phrase the DoA based decision as the simple vs. simple hypothesis test between H_0 ("all satellites nominal") and H_1 ("all satellites spoofed"). Subsets of spoofed satellites and analogously attacks from multiple, spatially distributed transmitting antennas can then be identified using the iteration presented in [8].

At any epoch t, a vector of DoA measurements $\boldsymbol{y_t}$ is, depending on the hypothesis, a function of the vector of true DoAs relative to the antenna's attitude $\boldsymbol{\phi_t}$, or the relative direction of spoofed signals $\boldsymbol{\phi_S}$ and some measurement noise. Measurements from k epochs $t_1, \ldots, t_k$ are concatenated in the vector $\boldsymbol{y_{t_1:t_k}}$. We generally model the measurement noise as zero-mean Normal and describe the calculation of its covariance in detail in the following sections for the cases of independent and correlated sequential measurements.

Prior work focuses on finding and working with a Maximum Likelihood Estimate (MLE) of the unknown antenna attitude and spoofer's direction at each epoch [2], [3], [5]. This comes at a significant computational cost and results in a less powerful hypothesis test as outlined for example by [2]. Instead, we follow an idea developed e.g. in [22], [23] of phrasing hypothesis independent of unknown quantities.

At time t, DoAs to N_t satellites in view generally deliver $2N_t$ measurements (e.g. N_t azimuth, N_t elevation). The unknown antenna's attitude contains three degrees of freedom, reducing the number of equations available for spoof detection to $2N_t - 3$. At each epoch, we therefore phrase $2N_t - 3$ adjusted measurement equations of great circle arcs between the satellite DoAs as they are independent of the antenna's orientation [24].
We will treat this general case of an unknown attitude for independent measurements in section III. For static antennas or applications with limited dynamics, small angles of pitch and bank can be assumed. The formulation of then $2N_t - 1$ adjusted measurements simplifies to $N_t - 1$ differences of azimuth and N_t elevation measurements. We derive the associated equations for the general case of potentially correlated measurements in section IV.

We formalize the adjusted measurements $\boldsymbol{\bar{y}_{t_1:t_k}}$ under either hypothesis in Equation (5). Under H_0 we expect the adjusted measurements to represent $2N_t - 3$ great circle arcs per epoch as calculated form the true azimuth and elevation values obtained from the ephemeris and user position and concatenated in $\boldsymbol{\bar{\phi}_{t_1:t_k}}$. Under H_1 the expected arcs are zero. The definition of the adjusted covariance matrix R depends on the precise definition of the noise on each DoA measurement and is covered separately in the two following sections.

$$\begin{aligned} H_0 &: \boldsymbol{\bar{y}_{t_1:t_k}} = \boldsymbol{\bar{\phi}_{t_1:t_k}} + \boldsymbol{\bar{\epsilon}_{t_1:t_k}} \\ H_1 &: \boldsymbol{\bar{y}_{t_1:t_k}} = \boldsymbol{\bar{\epsilon}_{t_1:t_k}} \end{aligned} \quad \text{with} \quad \boldsymbol{\bar{\epsilon}_{t_1:t_k}} \sim N(0, \bar{\mathrm{R}}) \tag{5}$$

Based on these measurement equations, the conditional probabilities $p\left(\boldsymbol{\bar{y}_{t_1:t_k}} | H_0\right)$ and $p\left(\boldsymbol{\bar{y}_{t_1:t_k}} | H_1\right)$ are evaluations of multivariate Normal distributions. Under H_0 with mean $\boldsymbol{\mu_0} = \boldsymbol{\bar{\phi}_{t_1:t_k}}$, under H_1 with zero mean $\boldsymbol{\mu_1} = 0$ and in either case with covariance matrix $\bar{\mathrm{R}}$. The natural log of the resulting likelihood ratio $\Lambda\left(\mathrm{y}_{\boldsymbol{t_1:t_k}}\right)$ develops into the result of Equation ().

$$
\begin{aligned}
\log\Lambda\left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right) = \log\frac{p\left(\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}|H_0\right)}{p\left(\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}|H_1\right)} &= -\frac{1}{2}\left(\left(\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}-\boldsymbol{\mu_0}\right)^T\bar{\mathrm{R}}^{-1}\left(\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}-\boldsymbol{\mu_0}\right)-\left(\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}-\boldsymbol{\mu_1}\right)^T\bar{\mathrm{R}}^{-1}\left(\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}-\boldsymbol{\mu_1}\right)\right)\\
&= \left(\boldsymbol{\mu_0}-\boldsymbol{\mu_1}\right)^T\bar{\mathrm{R}}^{-1}\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}-\frac{1}{2}\left(\boldsymbol{\mu_0}^T\bar{\mathrm{R}}^{-1}\boldsymbol{\mu_0}-\boldsymbol{\mu_1}^T\bar{\mathrm{R}}^{-1}\boldsymbol{\mu_1}\right)\\
&= \bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}\bar{\mathrm{R}}^{-1}\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}-\frac{1}{2}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}
\end{aligned}
\tag{6}
$$

The RV $\log\Lambda\left(\mathrm{y}_{\boldsymbol{t_1:t_k}}\right)$ is a linear function of the normally distributed variable $\boldsymbol{y}_{\boldsymbol{t_1:t_k}}$. We can compactly describe its expected distribution under nominal and spoofed conditions:

$$
\begin{aligned}
\log\Lambda\left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right)|H_0 &\sim N\left(\frac{1}{2}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}_{t_1:t_k}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}},\ \bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}_{t_1:t_k}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}\right)\\
\log\Lambda\left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right)|H_1 &\sim N\left(-\frac{1}{2}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}_{t_1:t_k}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}},\ \bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}_{t_1:t_k}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}\right)
\end{aligned}
\tag{7}
$$

The constraint of Eq. (4) and thereby the original optimization problem formulated in Eq. (1) is easily solved using the result in Eq. (7). The natural log of the detection threshold γ is given by Equation (8), where Φ^{-1} is the quartile function or inverse cdf of the Standard Normal distribution.

$$
\log\gamma = \frac{1}{2}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}_{t_1:t_k}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}} + \Phi^{-1}\left(P_{FA_{max}}\right)\sqrt{\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}_{t_1:t_k}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}}
\tag{8}
$$

We can calculate the detection probability or test power β for a spoofing attack where a set of malicious satellite signals is radiated from the same direction for a given satellite geometry, measurement accuracy and maximum false alert probability leveraging Equations (7) and (8) as a function of the cdf of the Standard Normal Distribution Φ.

$$
\beta = \Phi\left(\frac{\log\gamma + \frac{1}{2}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}_{t_1:t_k}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}\sqrt{\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}_{t_1:t_k}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}}}{\sqrt{\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}_{t_1:t_k}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}}}\right) = \Phi\left(\sqrt{\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{\mathrm{R}}_{t_1:t_k}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}} + \Phi^{-1}\left(P_{FA_{max}}\right)\right)
\tag{9}
$$

Equation (9) underlines the direct dependency of DoA based detection techniques on satellite geometry and measurement accuracy. A larger Mahalanobis distance $\frac{1}{2}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}^T\bar{R}^{-1}\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}$ results in a more powerful test.
This can serve as a design tool for the spoofing defense. For a given desired probability of detection β, maximum false alert probability $P_{FA_{max}}$ and most challenging geometry $\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}$ the designer can directly calculate the required measurement accuracy to be delivered by the hardware.

III. APPLICATION TO INDEPENDENT MEASUREMENTS

If different metrics like pseudorange residuals and DoAs are collected in $\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}$, they can likely be considered independent. In this section we expand the LRT for this case of independent measurements. We then develop the equations specifically for DoA measurements that are conditionally independent given the hypothesis and satellite geometry. We will then underline the general applicability of the procedure by reproducing previously published simulations.

1. Theoretical derivations

In the case of independent measurement noise, the computation of $\log\Lambda(\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}})$ in Eq. (4) simplifies to summing the values for each individual epoch as in Eq. (11). We show in the Appendix that the same result is obtained when modeling the state of being spoofed through a simple Hidden Markov Model (HMM) and calculating the posterior probability using a Histogram Filter [20].

$$
\log\Lambda\left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right) = \sum_{\tau=t_1}^{t_k}\log\frac{p\left(\bar{\boldsymbol{y}}_{\boldsymbol{\tau}}|H_0\right)}{p\left(\bar{\boldsymbol{y}}_{\boldsymbol{\tau}}|H_1\right)} = \sum_{\tau=t_1}^{t_k}\log\Lambda\left(\bar{\boldsymbol{y}}_{\boldsymbol{\tau}}\right) \quad \text{if} \quad Y_{t_1}\perp Y_{t_2}\ \forall\, t_1\neq t_2
\tag{10}
$$

This is a convenient result. $\log \Lambda \left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right)$ is the sum of the considered $\log \Lambda$s, each computed using Eq. (6). A sequential monitor can be implemented at practically no extra cost compared to a snapshot-based approach.

Analogous to the general formulations in Equations (7) and (9), the detection threshold $\log \gamma$ is given by Equation (12) and the test's power by Equation (13). Instead of stacking measurements in $\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}$ and $\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}$, we sum the Mahalanobis distances of each individual epoch for a more efficient computation.

$$\log \gamma = \frac{1}{2} \sum_{\tau=t_1}^{t_k} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}^T \bar{R}_\tau^{\,-1} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}} + \Phi^{-1}\left(P_{FA_{max}}\right) \sqrt{\sum_{\tau=t_1}^{t_k} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}^T \bar{\mathrm{R}}_\tau^{-1} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}} \tag{11}$$

$$\beta = \Phi\left(\frac{\log \gamma + \frac{1}{2}\sum_{\tau=t_1}^{t_k} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}^T \bar{\mathrm{R}}_\tau^{-1} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}} \sqrt{\sum_{\tau=t_1}^{t_k} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}^T \bar{\mathrm{R}}_\tau^{-1} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}}}{\sqrt{\sum_{\tau=t_1}^{t_k} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}^T \bar{\mathrm{R}}_\tau^{-1} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}}}\right) = \Phi\left(\sqrt{\sum_{\tau=t_1}^{t_k} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}^T \bar{\mathrm{R}}_\tau^{-1} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}} + \Phi^{-1}\left(P_{FA_{max}}\right)\right) \tag{12}$$

We note the low computational complexity of the computation of the decision threshold in Eq. (11). $\bar{R}$ is generally a banded matrix that can be efficiently inverted using a Cholesky factorization [25]. The size of involved matrices is naturally bounded by the total number of measurements for a given constellation. $\Phi^{-1}\left(P_{FA_{max}}\right)$ can be precomputed offline for a chosen false alert probability and stored in form of a lookup table for selected values $P_{FA_{max}}$.

The covariance of the adjusted measurements at epoch τ, $\bar{R}_\tau$ depends on the specific sensor setup and characterization. We model each DoA measurement noise as a rotation with Normally distributed magnitude in an arbitrary direction as done in [19], [26]. The distribution of the great circle arc between two such DoAs is nontrivial to describe. We model the noise of the arc between satellites i and j as zero-mean Normal with Variance $\bar{\sigma}^2 = \sigma_i^2 + \sigma_j^2$. Monte Carlo simulations of over 1 million arcs between randomized satellite positions with DoA measurement standard deviations between 3 deg and 15 deg have shown that the model over bounds the error, despite being a poor representation of small arcs due to their non-negativity. The interested reader is referred to the publicly available code underlying the simulations shown in the next subsection for an implementation example of the noise model.

2. Simulation results

The derivations in the previous subsections are valid for any approach that works with DoA measurements as obtained from e.g. a multi-element antenna array like the ones used by the groups that published [5], [6], as long as the measurement noise can be considered independent. To underline this applicability, we reproduce simulations presented in [5], [27] to the best of our knowledge as we have done in [15], here using three independent sequential measurements. We specifically show results based on the simulations underlying Figures 10 and 3 of both references [5], [27], which in this text we call scenario I and scenario II. We create two simulation setups with a fixed but randomly chosen constellation of 4 satellites and collect DoA measurements during 2000 and 3000 epochs respectively. During epochs [200, 400], [650, 800], [1000, 1500] in scenario I and [200, 400], [650, 800], [1000, 2000] in scenario II a spoofer is simulated. During these intervals, all signals are coming from the same azimuth and elevation before noise is added. Noise is generally modeled as independent and identically distributed (i.i.d.) for each satellite. It is realized as a rotation of the measurement vector in an arbitrary direction by a Normally distributed magnitude with variance of $15\,\mathrm{deg}^2$ in scenario I and $25\,\mathrm{deg}$ in scenario II. We show skyplots of each scenario in Figures 3a and 3b. The true satellite positions are depicted in red, the measured DoAs in blue. The spoofer's direction is simulated to be at around $115\,\mathrm{deg}$ azimuth and $46\,\mathrm{deg}$ elevation, corresponding to the cluster of measured DoAs in that direction in both skyplots.

In Figure 4 we show $\log \Lambda \left(\bar{\boldsymbol{y}}_{\boldsymbol{t-2:t}}\right)$ when normalized using its distribution under nominal conditions given by Eq. (7) for either scenario in green and blue, respectively. The subscript $t-2:t$ indicates that at each epoch t, we are using three sequential measurements at times $t-2$, $t-1$ and t. Dotted lines represent the snapshot based result as a baseline. Whenever the normalized value drops below the black line marking $\Phi\left(P_{FA_{max}}\right)$, an alarm is raised. In blue we indicate the spoofing flag. It is at the bottom of the plot for no alarm and at the top of the plot if an alarm is raised.

A decision based on $\log \Lambda \left(\bar{\boldsymbol{y}}_{\boldsymbol{t-2:t}}\right)$ is possible in either case with a large margin while guaranteeing a chosen false alert probability. The values of the normalized $\log \Lambda \left(\bar{\boldsymbol{y}}_{\boldsymbol{t-2:t}}\right)$ are significantly lower than the snapshot-based LRT depicted as dotted lines that we have presented on in [15].

The satellite geometry used in these simulations is likely different than in the cited literature, as the original geometry could not be determined. As we emphasized the test's power is entirely dependent on $\sum_{\tau=t_1}^{t_k} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}^T \bar{\mathrm{R}}_\tau^{-1} \bar{\boldsymbol{\phi}}_{\boldsymbol{\tau}}$ and and thereby the satellite geometry. Spoofing metrics are hence best compared on the exact same signal geometry and measurement noise. For this purpose and the benefit of the interested reader, the code used to generate these results is publicly available at `https://github.com/stanford-gps-lab/spoofing-detection`.

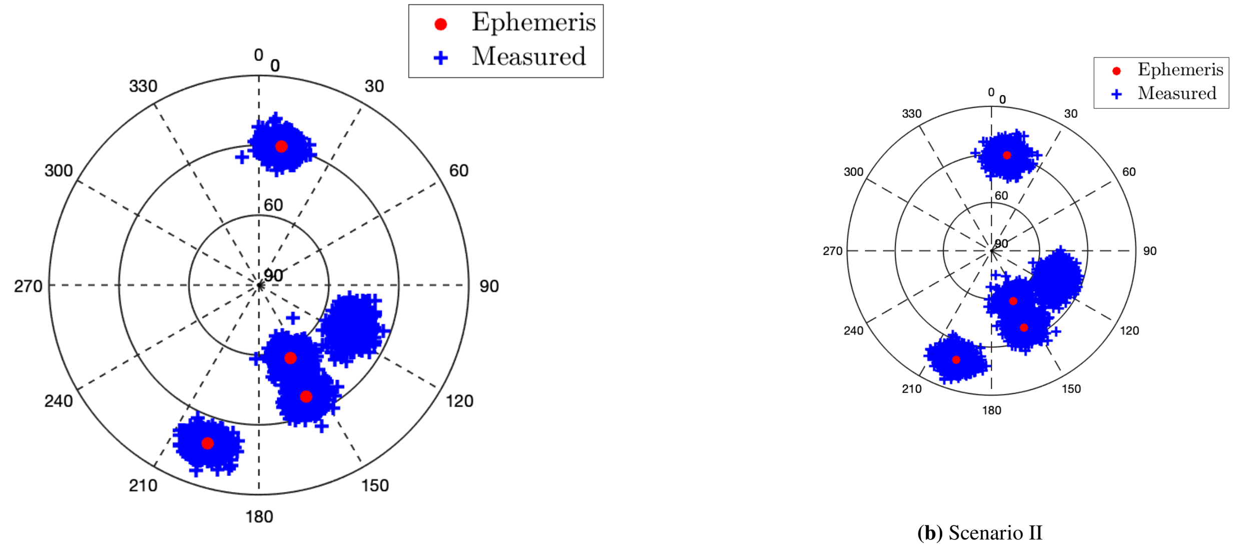

(a) Scenario I

(b) Scenario II

Figure 3: Skyplots of satellite positions based on ephemeris and DoA measurements. The spoofer's direction is simulated to be at around $115\,\mathrm{deg}$ azimuth and $46\,\mathrm{deg}$ elevation. The measurement variances are $15\,\mathrm{deg}^2$ in scenario I and $25\,\mathrm{deg}^2$ in scenario II.

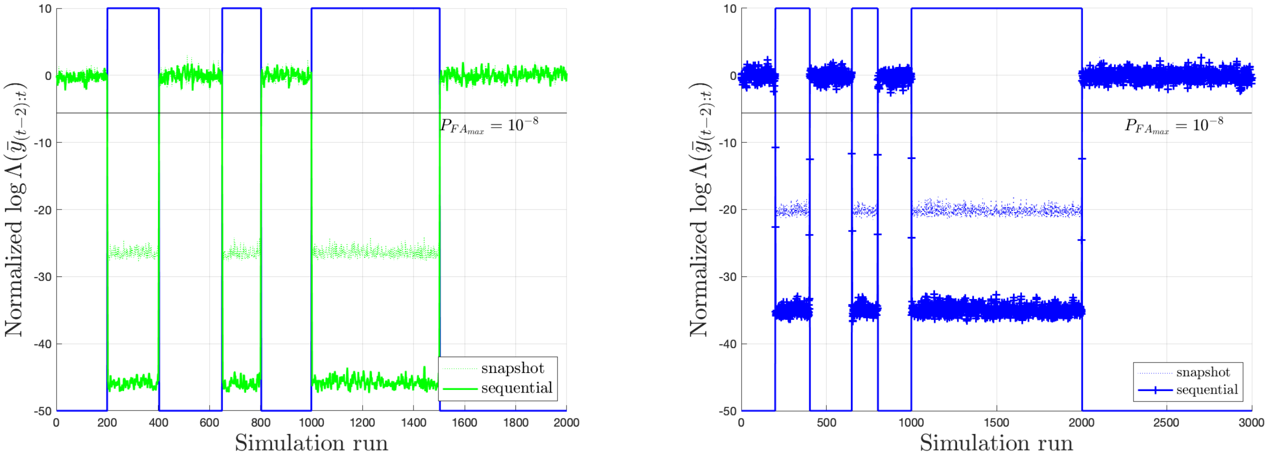

(a) Scenario I

(b) Scenario II

Figure 4: $\log \Lambda\left(\bar{\boldsymbol{y}}_{\boldsymbol{t-2:t}}\right)$ nomalized by its distribution under nominal conditions. The alert threshold for $P_{FA_{max}} = 10^{-8}$ is depicted as a black horizontal line.

IV. APPLICATION TO CORRELATED AZIMUTH MEASUREMENTS

We have now fully described adjusted measurements and a LRT using sequential DoA measurements that is optimal according to the Neyman Pearson lemma as long as the measurements are independent. For the remainder of this paper we will work with azimuth-only DoA measurements that can be obtained from a Dual Polarization Antenna (DPA) as we have shown in [7], [28], [29] and that have a timely correlation. In this subsection we illustrate how to account for the correlation in the LRT by constructing an appropriate covariance matrix. We validate the models performance through Monte Carlo Simulations. We show flight test data to characterize the correlation coefficient for the application example of an aircraft. We then apply the LRT to demonstrate the compliance with the false alert probability constraint. We finally apply the sequential LRT with correlation to data recorded with a DPA during a government sponsored live spoofing event to demonstrate the increased detection performance.

1. Theoretical Derivations

In this subsection we derive the adjusted sequential measurement $\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}$, truth $\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}$ and covariance $\bar{\mathrm{R}}_{t_1:t_k}$ for correlated azimuth measurements.

We show in [8] how the measurement covariance matrices at time t, R_t, can be estimated for a vector of azimuth measurements $\boldsymbol{y_t}$ from the DPA. The correlation between measurements at epochs t_1 and t_2 is described by the correlation coefficient $\rho_{t_1 t_2}$, which we characterize for an aircraft in flight in the next subsection.

With a characterized correlation coefficient we build the autocorrelation matrix $\mathrm{R}_{t_1:t_k}$ of the measurement vector $\boldsymbol{y}_{\boldsymbol{t_1:t_k}}$ for k correlated epochs t_1 through t_k as defined by Equation (13). This procedure is well documented in the literature, for example in [30]. The subscript ij indicates the matrix entry in the ith row and jth column. The covariance matrices of each individual epoch are symmetric and so is therefore the joint covariance matrix.

$$\mathrm{R}_{t_1:t_k} = \begin{bmatrix} \mathrm{R}_{t_1} & \mathrm{R}_{t_1 t_2} & \dots & \mathrm{R}_{t_1 t_{k-1}} & \mathrm{R}_{t_1 t_k} \\ \mathrm{R}_{t_2 t_1} & \mathrm{R}_{t_2} & & & \\ \vdots & & \ddots & & \vdots \\ \mathrm{R}_{t_{k-1} t_1} & & & \mathrm{R}_{t_{k-1}} & \\ \mathrm{R}_{t_k t_1} & & \dots & & \mathrm{R}_{t_k} \end{bmatrix} \quad \text{with} \quad \begin{matrix} \left(\mathrm{R}_{t_1 t_2}\right)_{ij} = \rho_{t_1 t_2} \sqrt{\left(\mathrm{R}_{t_1}\right)_{ij} \left(\mathrm{R}_{t_2}\right)_{ij}} \\ \mathrm{R}_{t_1 t_2} = \mathrm{R}_{t_1 t_2}^T = \mathrm{R}_{t_2 t_1} \end{matrix} \tag{13}$$

Since we are working with azimuth-only measurements, we assume small angles for the antenna's pitch and bank. When assuming small pitch and bank, only the unknown antenna heading remains to be eliminated. For a single epoch, the adjusted measurement and truth vectors $\bar{\boldsymbol{y}}_{\boldsymbol{t}}$ and $\bar{\boldsymbol{\phi}}_{\boldsymbol{t}}$ have therefore size $N_1 - 1$ and contain the differences of N_t DoA measurements and true azimuths as defined in Equation (14) using the sparse, banded reduction-matrix A_t.

$$\begin{matrix} \bar{\boldsymbol{y}}_{\boldsymbol{t}} = \mathrm{A}_t \, \boldsymbol{y_t} \\ \bar{\boldsymbol{\phi}}_{\boldsymbol{t}} = \mathrm{A}_t \, \boldsymbol{\phi_t} \end{matrix} \quad \text{with} \quad \mathrm{A}_t = \begin{bmatrix} 1 & -1 & \dots & 0 & 0 \\ 0 & 1 & -1 & \dots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & \dots & 1 & -1 \end{bmatrix} \in \mathbb{R}^{N-1 \times N} \tag{14}$$

The vectors $\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}$ and $\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}$ of k sequential, correlated epochs are constructed similarly. The sequential reduction matrix $\mathrm{A}_{t_1:t_k}$ is block diagonal constructed from the individual A_t. We summarize the definition of the adjusted sequential measurement, truth and covariance with correlation in Equation (15).

$$\begin{matrix} \bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}} = \mathrm{A}_{t_1:t_k} \, \boldsymbol{y}_{\boldsymbol{t_1:t_k}} \\ \bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}} = \mathrm{A}_{t_1:t_k} \, \boldsymbol{\phi}_{\boldsymbol{t_1:t_k}} \\ \bar{\mathrm{R}}_{t_1:t_k} = \mathrm{A}_{t_1:t_k} \, \mathrm{R}_{t_1:t_k} \, \mathrm{A}_{t_1:t_k}^T \end{matrix} \quad \text{with} \quad \mathrm{A}_{t_1:t_k} = \begin{bmatrix} \mathrm{A}_{t_1} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \mathrm{A}_{t_k} \end{bmatrix} \tag{15}$$

This result can be directly used in Equations (10) through (9) to calculate $\log \Lambda \left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right)$, its distribution under either hypothesis and the detection threshold and missed detection probability.

It is likely that DoA based spoof detection is not exclusively applied to all satellites in view of a constellation, for example to detect subsets of spoofed satellites as with the iterative algorithm presented in [15]. For that it is necessary to exclude satellites from the consideration. This can be done easily by adjusting the respective A_t matrices to have fewer rows and zeros in the respective columns.

Taking the correlation of measurements into account comes at a computational cost. The matrix $\mathrm{R}_{t_1:t_k}$ is no longer sparse or banded and grows in size as more measurements are considered, making its inversion computationally expensive. Limited computational resources in a receiver can be a second aspect besides the maximum TTA that limits the number of sequential measurements to be considered at once.

The timely correlation of measurements and respective non-sparsity of $\mathrm{R}_{t_1:t_k}$ creates a second point of concern. As we mentioned in [15] for snapshot based spoof detection based on DoA measurements, care has to be taken when wrapping the angular measurements.

The problem formulation in Eq. (1) constrains the maximum false alert probability. We therefore wrap each angular measurement such that it is closest to its expected value under nominal conditions for a conservative behavior towards the constraint, accepting

a possibly reduced detection power. Reference [23] goes into detail about the issue for values in the range of $[0,\ 1]$.
Due to the timely correlation in the sequential approach, a rare outlier case is possible. Here, the angular values in $\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}$ and $\bar{\boldsymbol{\phi}}_{\boldsymbol{t_1:t_k}}$ are generally in the range $[0,\ \pi]$. Say, for simplicity, the expected value of a measurement is 0 in two consecutive epochs with variance $\mathrm{R}_{t_1} = \mathrm{R}_{t_2} = (0.2\pi)^2$ and correlation coefficient $\rho_{t_1 t_2} = 0.5$. We actually measure angles of 0.4π and 0.6π, and the wrapping logic turns these into values of 0.4π and -0.4π (since -0.4π is closer to the expected value of 0 than 0.6π). Due to the correlation between measurements, this will result in an unreasonable low value of $p\left([0.4\pi\ (-0.4\pi)]^T|H_0\right) = 0.0015$, whereas it should have been closer to $p\left([0.4\pi\ (-0.4\pi)]^T|H_0\right) = 0.0015$, whereas it should have been closer to $p([0.4\pi\ 0.6\pi]^T|H_0) = 0.043$.

Instead of developing a complex wrapping logic trying to cover all outlier cases while still guaranteeing a the false alert limit, we simply choose the larger value between the calculated $p(\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}|H_0)$ using the wrapping logic and the product of conditional probabilities of the individual epochs. In our example, this would have been $p(0.4\pi|H_0)p(-0.4\pi|H_0) = 0.073$.

We validate our approach of sequential spoofing detection for correlated azimuthal DoA measurements with Monte Carlo simulations for up to four sequential epochs. All simulations are run for different, randomized geometries of 6 satellites. Measurement uncertainties are fixed at values between 17 deg and 29 deg and the correlation coefficients are $[0.62, 0.55, 0.5]^T$ for epochs that are 1, 2 and 3 time steps apart. Figure 5 shows the resulting values of $\Lambda\left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right)$, normalized normalized by its mean and standard deviation under nominal conditions, μ_{H_0} and σ_{H_0}, given by Eq. (7). Blue histograms show the result of 10^5 simulations of nominal conditions per number of epochs, red histograms show the result of 10^5 simulations of spoofed conditions per number of epochs. The black lines show standard Normal distributions, representing the expected distribution of the blue histograms. An alarm would be raised for any value below $\Phi^{-1}\left(P_{FA_{max}}\right)$. We can see that the model represented by the black $P_{FA_{max}}$ standard Normal distributions is conservative, the blue histograms show a shift towards positive values due to the described conservative computation of $p(\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}|H_0)$ for sequential measurements. We further note how a larger number of sequential measurements causes a reduction of the red $\left(\Lambda\left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right) - \mu_{H_0}\right)/\sigma_{H_0}$ values under spoofed conditions, indicating more powerful tests.
To quantify the gain in detection performance, we plot the receiver operating characteristics (ROCs) in Figure 6. Specifically, we show the empirical probability of missed detection (P_{MD}) plotted against the maximum false alert probability for up to four sequential epochs. Note the log scale for $P_{FA_{max}}$. The empirical P_{MD} is the percentage of values of the spoofed $\left(\Lambda\left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right) - \mu_{H_0}\right)/\sigma_{H_0} < \Phi^{-1}\left(P_{FA_{max}}\right)$ in the data shown in Figure 5. Each line can be considered a Pareto frontier in the trade-off between missed detections and false alerts [31]. We observe a significant performance improvement, for constant P_{MD} we e.g. achieve a reduction of $P_{FA_{max}}$ by several orders of magnitude when using the sequential approach. Since the data was generated from randomized geometries, we consider this performance gain a representative average.

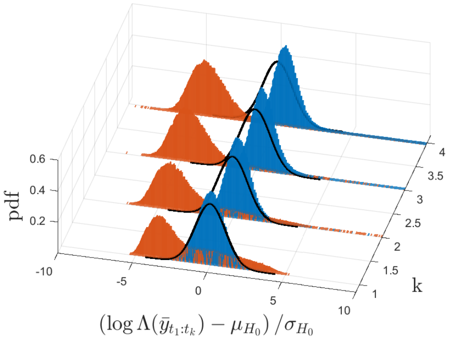

Figure 5: Comparison of $\Lambda\left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right)$ when based on the measurements from 1 through 4 epochs. 10^5 measurements from 6 satellites per epoch for both nominal (blue) and spoofed (red) conditions.

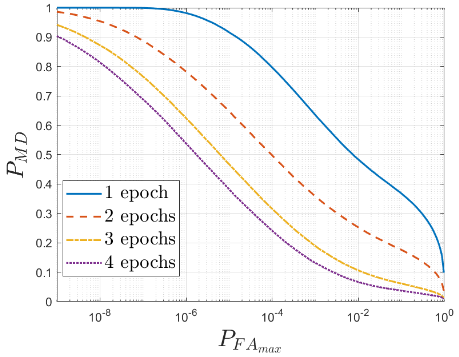

Figure 6: ROCs when using measurements from 1 through 4 epochs based on the data shown in Figure 5. P_{MD} is plotted on a log scale.

2. Application to an aircraft in flight

In this subsection we apply the theoretical derivations to flight data recorded by a DPA mounted on a Beechcraft C-12 aircraft. This was a series of test conducted in September 2019 and is described in [29]. We characterize the correlation coefficient and show the distribution of $\log\Lambda(\bar{\boldsymbol{y}})$ using a snapshot and sequential approach to further validate the model and demonstrate the compliance with the maximum false alert probability.

Small pitch and especially bank angles is not always a valid assumption for an aircraft. Attitude changes however only happen in flight, requiring a significant velocity. The GPS course over ground vector can then be leveraged for a reasonable attitude estimate when assuming coordinated turns, which is again a reasonable assumption for most civil aircraft. We inflate the covariance as a function of the bank angle, leading to a realistic estimate of R throughout the flight as we have shown in [15]. If a more precise attitude estimate is desired, the DoA measurement could be coupled with MEMS gyros.

The correlation coefficient $\rho_{t_1 t_2}$ of Eq. (13) can be determined by calculating the autocorrelation of the measurement errors [32]. In Figure 7 we show the autocorrelation of the azimuth measurement error as a function of time delay $|t_1 - t_2|$. The data was recorded on the C-12 aircraft during a flight profile of various climbs, descends and turns with up to 60 deg bank. Of course, $\rho_{t_1 t_2}$ is highly dependent on the application and situation. A car driving in an urban environment for example has a different correlation than a well calibrated DoA measuring antenna on the roof on an aircraft in flight. The slower movements and stronger multipath would likely result in a stronger correlation. For our application example and based on the results depicted in Figure 7 we set the correlation coefficient values as shown in Table 1.

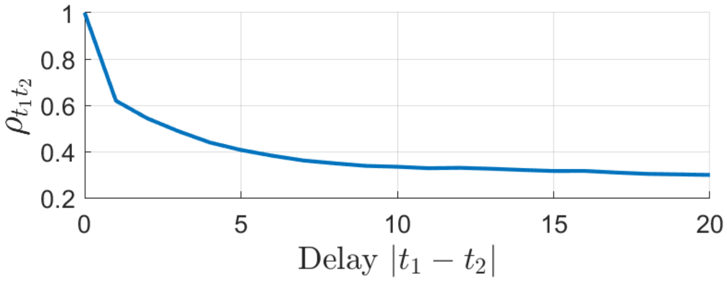

Figure 7: Autocorrelation of the azimuth measurement error recorded by the DPA in flight

Table 1: Correlation values determined for the DPA on an aircraft in flight

Delay $\|t_1 - t_2\|$	1	2	3
Autocorrelation $\rho_{t_1 t_2}$	0.62	0.55	0.5

To validate the derived theoretical model of the measurements and correlation, we calculate $\log \Lambda(\bar{\boldsymbol{y}})$ with both a snapshot-based approach of considering individual measurements as well as the presented sequential approach with two consecutive measurements using Eq. (6). We normalize each by its expected distribution under nominal conditions given by Eq. (7). To match the model, the resulting variable should resemble a Standard Normal distribution. We show histograms of the resulting normalized $\log \Lambda(\bar{\boldsymbol{y}})$ values collected during a 2.5h flight in Figure 8 and Figure 9 for the snapshot and sequential approach, respectively, plotted against a Standard Normal probability density function (pdf).

Both distributions overall resemble and are over bound by the Standard Normal distribution. The bounding is caused by the measurement covariance R_t computed as presented in [8] which over bounds the measurement error. The measurements are more precise than expected, resulting in a tighter distribution of $\log \Lambda(\bar{\boldsymbol{y}})$. An alarm is raised if the normalized $\log \Lambda(\bar{\boldsymbol{y}}) < \Phi^{-1}\left(P_{FA_{max}}\right)$, the shown distributions therefore result in a conservative behavior towards the false alert probability constraint. The distribution of the sequential $\log \Lambda\left(\bar{\boldsymbol{y}}_{\boldsymbol{t-2:t}}\right)$ closely resembles that of the snapshot-based result, with no large outliers towards the negative side that would have jeopardized the false alert probability constraint. The scope of this paper is not a comprehensive characterization of this specific DPA mounted on a Beechcraft C-12, but rather a first feasibility study. For this purpose, we consider this result sufficient to validate the correlation model.

3. Application to live spoofing data

After confirming the behavior under nominal conditions, we apply the sequential LRT to data collected during a live spoofing event sponsored by the US government that we have previously presented on in [7], [8], [15].

Data was recorded during a total of 39 episodes of spoofing, always from a single source transmitter. Each episode was fairly short compared to the measurement frequency, resulting in 1-7 measurement epochs per spoofed episode. During most spoofed epochs we received a mix of genuine and spoofed signals. Each constellation was processed separately, for a total of 429 measurement epochs of which 124 were spoofed during the 39 episodes of consecutive spoofing.
The antenna was mounted on top of a car using a roof rack in an environment that resembles a half full parking lot. Without the opportunity to extensively characterize the antenna behavior and measurement correlation in this environment, we use the correlation values derived in the previous subsection for an aircraft in flight.

In Figure 10 we show the percentage of expected and actual detections using measurements from 1 to 4 sequential epochs for

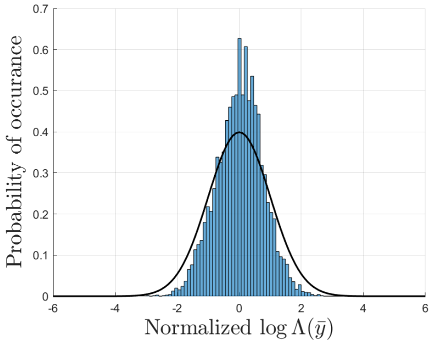

Figure 8: Histogram of normalized snapshot $\log\Lambda(\bar{\boldsymbol{y}})$ values during a 2.5h flight. The black line is a Standard Normal pdf and resembles the measurement model

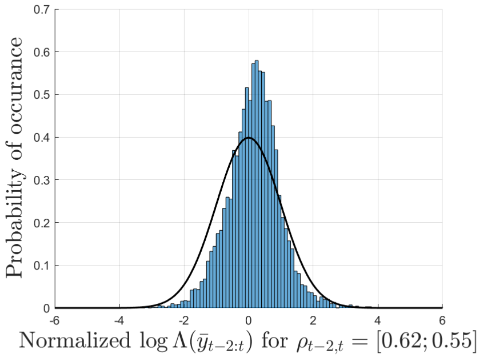

Figure 9: Histogram of normalized sequential $\log\Lambda(\bar{\boldsymbol{y}}_{\boldsymbol{t-2:t}})$ values during a 2.5h flight. The black line is a Standard Normal pdf and resembles the measurement model

a conservative maximum false alert probability of $P_{FA_{max}} = 10^{-7}$. The expected percentage is the average of the test power β computed using Eq. (9) for each test. In Table 2 we summarize the increase of both expected and actual detections of the sequential approach relative to the snapshot approach using a single epoch.

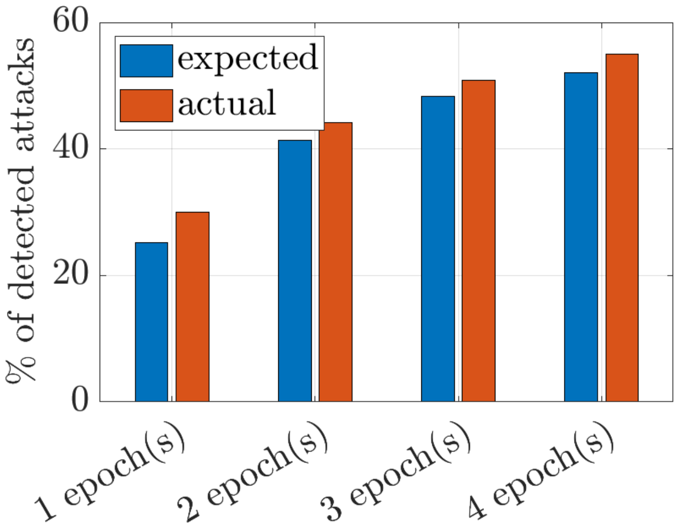

Figure 10: Theoretical power β and percentage of detected attacks using up to four sequential epochs during the government sponsored live spoofing event.

Figure 10 shows a continuous increase of theoretical power and actual percentage of detected attacks for an increasing number of sequential measurement epochs. Due to the very limited number of measurement epochs during each episode of consecutive spoofing, added epochs can only be leveraged at fewer instances as more epochs are considered. The added benefit of each additional epoch therefore decreases with each additional epoch, even more so than in the simulation results shown in Figure 5. From Table 2 we see that using 4 epochs doubles the theoretical test power and results in an increase of detections by 83%.
For each number of epochs more attacks are detected than forecasted according to the model. This is not unexpected. The spoofed signals are coming from the same direction and experience some common mode errors (multipath, antenna imperfections, ...). This causes their DoAs to be closer aligned with a better fit to $p\left(\bar{\boldsymbol{y}}_{\boldsymbol{t_1:t_k}}|H_1\right)$ than modeled, resulting in a stronger LRT. The model of H_1 in Eq. (7) is conservative.

Table 2: Correlation values determined for the DPA on an aircraft in flight

Number of epochs	2	3	4
Increase of av. power β	63.9%	91.5%	106.6%
Increase of detections	47.2%	69.4%	83.3%

V. SUMMARY AND CONCLUSIONS

In this paper we present a scheme to combine multiple measurements for GNSS spoof detection in a LRT with provable alert behavior for safety of life applications. We formulate an algorithm for combining both independent and correlated DoA measurements that result in an analytic solution for the detection threshold that can be computed online by the receiver. In Monte Carlo simulations the sequential approach achieves a reduction of the false alert probability by multiple orders of magnitude for constant detection power compared to snapshot based detection.

We show the general applicability of the scheme by reproducing previously published simulations. We then validate the scheme for correlated azimuth measurements with data recorded by a DPA mounted on a C12 aircraft in flight. We finally apply the procedure to data from a live spoofing event and show an increase in detections by 47% already using just two sequential measurements, with equal robustness to false alerts compared to snapshot based detection.

These results show that using sequential spoof detection is a powerful way to improve the detection capability of an anti-spoof defense. It comes at the cost of added computational complexity and introduces a timely component to the detection procedure that points towards a necessary discussion about a maximum TTA.

Moving forward the authors hope to expand the LRT scheme towards the use of multiple different metrics, for a statistically sound decision about the presence of a spoofing attack using all the information available to a receiver.

ACKNOWLEDGMENTS

The authors thank the Federal Aviation Administration (FAA) and the Stanford Center for Position Navigation and Time (SCPNT) for sponsoring this research. The authors also thank the US government for providing us with an opportunity to test under live GPS spoofing. Special thanks goes to the entire team of the Air Force Test Center at Edwards Air Force Base and the 586th flight test squadron for their support leading up to and the conduct of the flight tests. The authors thank Prof. J. David Powell for his invaluable advice and guidance during this research.

REFERENCES

[1] M. L. Psiaki and T. E. Humphreys, “GNSS Spoofing and Detection,” *Proceedings of the IEEE*, vol. 104, no. 6, pp. 1258–1270, 2016.

[2] M. L. Psiaki, S. P. Powell, and B. W. O’Hanlon, “GNSS Spoofing Detection Using High-Frequency Antenna Motion and Carrier-Phase Data,” in *26th International Technical Meeting of the Satellite Division of the Institute of Navigation*, 2013, pp. 2949–2991. [Online]. Available: https://gps.mae.cornell.edu/Paper{_}F5{_}8{_}ION{_}GNSS{_}2013b.pdf

[3] M. L. Psiaki, B. W. O’Hanlon, S. P. Powell, J. A. Bhatti, T. E. Humphreys, and A. Schofield, “GNSS Spoofing Detection using Two-Antenna Differential Carrier Phase,” in *Proceedings of the 27th International Technical Meeting of The Satellite Division of the Institute of Navigation*, 2014, pp. 2776–2800. [Online]. Available: http://hdl.handle.net/2152/63214

[4] G. Falco, M. Nicola, E. Falletti, and M. Pini, “An Algorithm for Finding the Direction of Arrival of Counterfeit GNSS Signals on a Civil Aircraft,” in *Proceedings of the 27th International Technical Meeting of the Satellite Division of The Institute of Navigation, ION GNSS+ 2019*, Miami, Florida, 2019, p. 12.

[5] M. Appel, A. Iliopoulos, F. Fohlmeister, E. Pérez Marcos, M. Cuntz, A. Konovaltsev, F. Antreich, and M. Meurer, “Experimental validation of GNSS repeater detection based on antenna arrays for maritime applications,” *CEAS Space Journal*, vol. 11, no. 1, pp. 7–19, 2019. [Online]. Available: http://dx.doi.org/10.1007/s12567-018-0232-6

[6] M. C. Esswein and M. L. Psiaki, “Vector tracking of GNSS signals from a multi-element antenna,” in *Proceedings of the 31st International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS+ 2018*, no. 1, 2018, pp. 4040–4051.

[7] Y.-H. Chen, S. Lo, A. Perkins, F. Rothmaier, D. M. Akos, and P. Enge, “Demonstrating Single Element Null Steering Antenna Direction Finding for Interference Detection,” *Proceedings of the 2018 International Technical Meeting of The Institute of Navigation*, pp. 240–259, 2018.

[8] F. Rothmaier, Y.-h. Chen, and S. Lo, "Improvements to Steady State Spoof Detection with Experimental Validation using a Dual Polarization Antenna," in *Proceedings of the 2019 International Technical Meeting of The Institute of Navigation ION GNSS+*, Miami, Florida, 2019.

[9] Dennis M. Akos, "Who's Afraid of the Spoofer? GPS/GNSS Spoofing Detection via Automatic Gain Control (AGC)," *NAVIGATION, Journal of the Institute of Navigation*, vol. 59, no. 4, pp. 281–290, 2012. [Online]. Available: https://www.ion.org/publications/abstract.cfm?articleID=102583

[10] E. G. Manfredini, D. M. Akos, Y.-H. Chen, S. Lo, T. Walter, and P. Enge, "Effective GPS Spoofing Detection Utilizing Metrics from Commercial Receivers," in *Proceedings of the 2018 International Technical Meeting of The Institute of Navigation*, 2018, pp. 672–689.

[11] C. Sun, J. W. Cheong, A. G. Dempster, H. Zhao, and W. Feng, "GNSS spoofing detection by means of signal quality monitoring (SQM) metric combinations," *IEEE Access*, vol. 6, pp. 66 428–66 441, 2018.

[12] A. Pirsiavash, A. Broumandan, and G. Lachapelle, "Characterization of signal quality monitoring techniques for multipath detection in GNSS applications," *Sensors (Switzerland)*, vol. 17, no. 7, 2017.

[13] C. Hegarty, B. O'Hanlon, A. Odeh, K. Shallberg, and J. Flake, "Spoofing detection in GNSS receivers through cross-ambiguity function monitoring," in *Proceedings of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS+ 2019*. Miami, Florida: Institute of Navigation, 2019, pp. 920–942.

[14] J. N. Gross, C. Kilic, and T. E. Humphreys, "Maximum-likelihood power-distortion monitoring for GNSS-Signal authentication," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, no. 1, pp. 469–475, 2019.

[15] F. Rothmaier, Y.-H. Chen, S. Lo, and T. Walter, "GNSS Spoof Detection through Spatial Processing," *Submitted for publication in NAVIGATION*, 2020.

[16] J. A. Rice, *Mathematical Statistics and Data Analysis*, 3rd ed. Belmont, CA: Thomson Brooks/Cole, 2007.

[17] P. Montgomery, T. E. Humphreys, and B. Ledvina, "Receiver-autonomous spoofing detection: Experimental results of a multi-antenna receiver defense against a portable civil GPS spoofer," in *Proceedings of the international Technical Meeting of the Institute of Navigation 2009*, vol. 1, Anaheim, CA, 2009, pp. 124–130.

[18] P. D. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*, 1st ed. Boston, NY: Artech House, 2008.

[19] M. Meurer, A. Konovaltsev, M. Appel, and M. Cuntz, "Direction-of-Arrival Assisted Sequential Spoofing Detection and Mitigation," in *Proceedings of the 2016 International Technical Meeting of The Institute of Navigation*, 2016, pp. 181–192.

[20] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. Cambridge, Massachusetts: The MIT Press, 2005.

[21] H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I*. New York: John Wiley & Sons, Inc., 2001.

[22] J. Magiera and R. Katulski, "Detection and mitigation of GPS spoofing based on antenna array processing," *Journal of Applied Research and Technology*, vol. 13, no. 1, pp. 45–57, 2015. [Online]. Available: http://dx.doi.org/10.1016/S1665-6423(15)30004-3

[23] D. Borio and C. Gioia, "A sum-of-squares approach to GNSS spoofing detection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 4, pp. 1756–1768, 2016.

[24] D. T. Greenwood, *Principles of Dynamics*. Englewood Cliffs, New Jersey: Prentice Hall, 2012.

[25] E. Kiliç and P. Stanica, "The inverse of banded matrices," *Journal of Computational and Applied Mathematics*, vol. 237, no. 1, pp. 126–135, 2013. [Online]. Available: http://dx.doi.org/10.1016/j.cam.2012.07.018

[26] M. Meurer, A. Konovaltsev, M. Cuntz, and C. Hättich, "Robust joint multi-Antenna spoofing detection and attitude estimation using Direction assisted multiple hypothese RAIM," in *Proceedings of the 25th Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2012)*, no. May 2016, 2012. [Online]. Available: https://elib.dlr.de/79200/1/Meuer{_}Robust{_}Joint{_}Multi-Antenna{_}Session{_}D5{_}final.pdf

[27] M. Appel, A. Konovaltsev, and M. Meurer, "Robust Spoofing Detection and Mitigation based on Direction of Arrival Estimation," in *Proc. ION GNSS+ 2015*, 2015, pp. 3335–3344. [Online]. Available: https://www.ion.org/gnss/abstracts.cfm?paperID=2911

[28] Y.-h. Chen, F. Rothmaier, D. M. Akos, S. Lo, and P. Enge, "PCB Implementation of a Single Null-steering Antenna and its Anti-spoofing / jamming Testing," *Proceedings of the 2017 International Technical Meeting of The Institute of Navigation*, 2017.

[29] S. Lo, Y.-H. Chen, F. Rothmaier, G. Zhang, and C. Lee, “Developing a Dual Polarization Antenna (DPA) for High Dynamic Applications,” in *Proceedings of the International Technical Meeting (ITM) of the Institute of Navigation (ION) 2020*, San Diego, CA, 2020.

[30] D. Simon, *Optimal State Estimation*, 1st ed. Hoboken, New Jersey: John Wiley & Sons, Inc., 2006. [Online]. Available: https://academic.csuohio.edu/simond/estimation/

[31] M. J. Kochenderfer and T. A. Wheeler, *Algorithms for Optimization*, 1st ed. Cambridge, Massachusetts: The MIT Press, 2018.

[32] G. E. Box, G. M. Jenkins, and G. C. Reinsel, *Time series analysis: Forecasting and control*, 4th ed. Englewood Cliffs, New Jersey: Prentice Hall, 2013.

APPENDIX

The result of Equation (10) can also be obtained by modeling the state of being spoofed as a Hidden Markov Model (HMM) and combining a Histogram Filter architecture with a statistical hypothesis test [20].

We start by replacing the prior probabilities of any epoch by the posterior probability of its preceding epoch multiplied with the transition probability matrix T as formalized in Equations (16) and (17).

$$\boldsymbol{\pi_t} = \mathrm{T}\left[p(H_0|\boldsymbol{y_{1:t-1}})\; p(H_1|\boldsymbol{y_{1:t-1}})\right]^T \tag{16}$$

$$\mathrm{T} = \begin{bmatrix} \mathrm{T}_{00} & \mathrm{T}_{01} \\ \mathrm{T}_{10} & \mathrm{T}_{11} \end{bmatrix} = \begin{bmatrix} p(x_t = H_0|x_{t-1} = H_0) & p(x_t = H_0|x_{t-1} = H_1) \\ p(x_t = H_1|x_{t-1} = H_0) & p(x_t = H_1|x_{t-1} = H_1) \end{bmatrix} \tag{17}$$

where

$$\sum_j \mathrm{T}_{ij} = 1 \;\forall\; i \tag{18}$$

and x_t is the state at time t. The relationship in Equations (16) and (17) contains the assumption inherent to an HMM: x_t is a Markov Blanket with respect to all preceding states and measurements [20]. The measurements $\boldsymbol{y_t}$ are conditionally independent given their respective state x_t, the measurement error ϵ_t is independent and identically distributed (i.i.d.).
Combining the definition of the prior probability in Equations (16) and (17) with Bayesian state estimation of a continuous variable is commonly known as a Histogram Filter [20]. Following the standard procedure, we calculate the ratio of prior probabilities as time t through the recursive expression in Eq. (19).

$$\frac{\pi_t(H_0)}{\pi_t(H_1)} = \frac{\mathrm{T}_{00}\, p(H_0|\boldsymbol{y_{1:t-1}}) + \mathrm{T}_{01}\, p(H_1|\boldsymbol{y_{1:t-1}})}{\mathrm{T}_{10}\, p(H_0|\boldsymbol{y_{1:t-1}}) + \mathrm{T}_{11}\, p(H_1|\boldsymbol{y_{1:t-1}})} \tag{19}$$

where $p(H_i|\boldsymbol{y_{1:t-1}}$ is given by Bayes Rule.

$$p(H_i|\boldsymbol{y_{1:t-1}} = \frac{p(\boldsymbol{y_{t-1}}|H_i)\pi_{t-1}(H_i)}{p(\boldsymbol{y_{1:t-1}})} \tag{20}$$

We can then rewrite Eq. (19)

$$\frac{\pi_t(H_0)}{\pi_t(H_1)} = \frac{\mathrm{T}_{00}\, p(\boldsymbol{y_{t-1}}|H_0)\pi_{t-1}(H_0) + \mathrm{T}_{01}\, p(\boldsymbol{y_{t-1}}|H_1)\pi_{t-1}(H_1)}{\mathrm{T}_{10}\, p(\boldsymbol{y_{t-1}}|H_0)\pi_{t-1}(H_0) + \mathrm{T}_{11}\, p(\boldsymbol{y_{t-1}}|H_1)\pi_{t-1}(H_1)} \tag{21}$$

We now make one simplification. The transition probabilities T_{ij} are chosen by the designer, expressing the belief of how the situation changes from one epoch to the next. The off-diagonal probabilities T_{01} and T_{10} are likely to be small, as the state can be expected to rarely change from nominal to spoofed and back. To simplify the following steps, we therefore set them to zero. With $\mathrm{T}_{01} = \mathrm{T}_{10} = 0$ Equation (21) simplifies:

$$\frac{\pi_t(H_0)}{\pi_t(H_1)} = \frac{p(\boldsymbol{y_{t-1}}|H_0)}{p(\boldsymbol{y_{t-1}}|H_1)} \frac{\pi_{t-1}(H_0)}{\pi_{t-1}(H_1)} \tag{22}$$

Applying Equation (22) recursively on all measurements between two epochs t_1 and t_{k-1}:

$$\frac{\pi_{t_k}(H_0)}{\pi_{t_k}(H_1)} = \prod_{\tau=t_1}^{t_{k-1}} \frac{p(\boldsymbol{y_\tau}|H_0)}{p(\boldsymbol{y_\tau}|H_1)} \frac{\pi_{t_1}(H_0)}{\pi_{t_1}(H_1)} \tag{23}$$

Substituting Eq. (23) into Eq. (3) and again collecting RVs on one side:

$$p\left(\prod_{\tau=t_1}^{t_k} \frac{p(\boldsymbol{y_\tau}|H_0)}{p(\boldsymbol{y_\tau}|H_1)} < \left(\frac{1}{c}-1\right) \frac{\pi_{t_1}(H_0)}{\pi_{t_1}(H_1)}|H_0\right) \leq P_{FA_{max}} \tag{24}$$

We collapse constants into the sequential threshold γ and take the natural log to arrive at the same result as Equation (10).

$$\log\Lambda\left(\bar{\mathrm{y}}_{\boldsymbol{t_1:t_k}}\right) = \sum_{\tau=t_1}^{t_k} \log\frac{p\left(\boldsymbol{y_\tau}|H_0\right)}{p\left(\boldsymbol{y_\tau}|H_1\right)} = \sum_{\tau=t_1}^{t_k} \log\Lambda\left(\boldsymbol{y_\tau}\right) \quad \text{if} \quad Y_{t_1} \perp Y_{t_2} \; \forall \; t_1 \neq t_2 \tag{25}$$