

GNSS spoofing detection through spatial processing

Fabian Rothmaier  | Yu-Hsuan Chen | Sherman Lo  | Todd Walter 

Aeronautics & Astronautics Department,
Stanford University, Stanford, California,
USA

Correspondence

Fabian Rothmaier, Aeronautics &
Astronautics Department, Stanford Uni-
versity, Stanford, CA 94305-4035 USA
Email: fabianr@stanford.edu

Funding information

Federal Aviation Administration; Memo-
randum MOA 693KA8-19-N-00015

Abstract

In this paper, we present an algorithmic framework for signal-geometry-based approaches of GNSS spoofing detection. We formulate a simple vs. simple hypothesis test independent of nuisance parameters that results in significantly reduced missed detection probability compared to prior approaches. It is highly tractable such that it can be computed online by the receiver. We employ a hypothesis iteration framework that finds spoofed subsets of satellites efficiently and accounts for the presence of weak multipath, for a provable decision behavior in safety-of-life applications. We support the theoretical derivations by showing results on previously published simulated and on-air data sets. We validate the measurement model and show robustness to multipath with flight data from a Dual Polarization Antenna (DPA) mounted on a C12 aircraft. Finally, we show the algorithm's benefit on data recorded during a government-sponsored live spoofing event.

KEYWORDS

direction of arrival, spoofing detection, statistical testing

1 | INTRODUCTION

With around as many GNSS receivers in the world as people with access to electricity, satellite navigation has become a ubiquitous technology that is constantly relied on (European Global Navigation Satellite Systems Agency, 2017). It is being used increasingly to support autonomy in applications such as drones, vessels, railway and autonomous cars. Given the reliance on GNSS in many autonomous applications, GNSS receivers will need to be able to provide high integrity in all environments – even in the presence of interference such as spoofing.

The vulnerability of current receivers to spoofing has been demonstrated, for example, in (Humphreys et al., 2008) and (Bhatti & Humphreys, 2017), and GNSS interfer-

ences are recognized as “serious threats to the continued safety of air transport” (RASG-MID, 2019).

Robust GNSS spoofing detection is a field of active research. Many possible detection means have been proposed; however, there is no single panacea. The goal is to make it much too costly or not worthwhile given the effort required for an attacker to attempt to overcome our implemented defense. In general, the better our defense, the more expensive it will be to mount a successful attack. References (Günther, 2014; Jafarnia-Jahromi et al., 2012; Psiaki & Humphreys, 2016) give overviews of current attack and defense strategies.

The proposed defenses have benefits and limitations. For example, monitoring automatic gain control (AGC)/input power (Akos, 2012), carrier-to-noise ratio (CN0)/

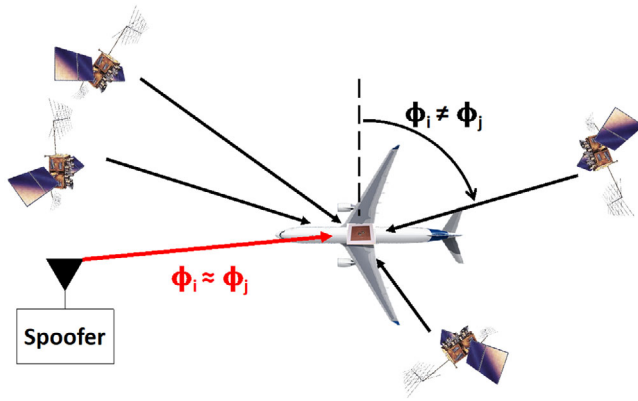


FIGURE 1 The signal-geometry-based concept from a bird's-eye perspective: Genuine signal directions (black) are diverse, while all spoofing signal directions (red) align [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

signal power and/or the correlation function (Gross et al., 2019; Manfredini et al., 2018; Pirsivash et al., 2016; Wesson et al., 2018) are easy to implement and require no hardware additions. A major limitation of these techniques is that they work only during the initial capture phase of an event. These so-called “transient detectors” cannot detect an attack once the spoofer has captured the receiver. If he succeeds in dragging the victim off the original correlation function unnoticed, for example, by jamming the victim's receiver first, the attack likely will continue unnoticed. While the detector presented in (Gross et al., 2019; Wesson et al., 2018) is designed to detect the jamming attack, it is unclear how it would differentiate between authentic and spoofed signals after liftoff.

It is useful to couple a “steady state detection,” whereby detection can occur at any time, not just during the capture phase. One steady state detection approach is to use the spatial diversity present in GNSS signals. Generally at the expense of hardware changes, like using a dithering antenna (Psiaki et al., 2013), two antennas (Borio & Gioia, 2016; Psiaki et al., 2014), an entire array of antennas (Appel et al., 2019; Appel et al., 2015; Esswein & Psiaki, 2019; Konovaltsev et al., 2014; Konovaltsev et al., 2013; Magiera & Katulski, 2015; Meurer et al., 2016; Meurer et al., 2012) or a Dual Polarization Antenna (DPA) (Lo et al., 2018; Lo et al., 2020), metrics reflecting the different directions of arrival (DoA)s of the GNSS signals are derived. Under nominal conditions, these metrics will be different for each satellite as an antenna receives signals from satellites distributed across the entire sky. Signals transmitted from a spoofer, on the other hand, will arrive from a single or a few directions, if an expensive attack using multiple transmitting antennas is mounted. Figure 1 shows a bird's-eye view of the concept for four satellites. Angles/directions of

signals i and j are different when coming from the authentic satellites (black arrows) but near identical when coming from a single spoofing source (red arrow).

This underlines the hypothesis inherent to signal-geometry-based approaches, that all – or at least multiple – satellite signals will be broadcasted from the same source by the attacker. The cited literature demonstrates strong results for situations when all GPS signals are malicious and transmitted from a single antenna and shows an analysis of false alert and missed detection rates for the respective cases. (Esswein & Psiaki, 2019) specifically further analyzes its performance in detecting spoofed subsets through an optimization-based approach.

This paper makes three contributions to spoofing detection based on measured DoAs.

First, we cast the detection as a hypothesis test that guarantees a chosen false alert probability that is the Uniformly Most Powerful Invariant (UMPI) test independent of nuisance parameters (Lehmann & Romano, 2005). We formulate hypotheses that enable a fast, online computation of the detection threshold for any DoA-based approach that has direction measurements at its disposal, like the techniques described by (Appel et al., 2019; Appel et al., 2015; Esswein & Psiaki, 2019; Konovaltsev et al., 2014; Konovaltsev et al., 2013; Meurer et al., 2016; Meurer et al., 2012). We compare our results to several existing approaches and achieve a more than 50% lower missed detection probability.

The second contribution is an updated version of a greedy hypothesis iteration algorithm that we presented first in (Rothmaier, Chen, & Lo, 2019). It efficiently and effectively identifies subsets of spoofed satellites, even in the presence of weak multipath.

As a third contribution, we show how these considerations are necessary for reasonable performance under real-world conditions. We present flight test data and data collected during a government-sponsored live spoofing event to support the theoretical derivations.

The paper covers these contributions in three main sections, plus a summary and conclusions. Section 2 sets up a Neyman-Pearson Likelihood Ratio Test (LRT) for a given maximum false alert probability. We detail the UMPI-based algorithm and apply it to several approaches presented in the literature. We show the direct dependence of any DoA-based approach on satellite geometry and DoA measurement accuracy and outline limitations of the approach.

In Section 3, we present the hypothesis iteration algorithm. It mitigates weak multipath and efficiently detects subsets of spoofed satellites. We show flight test data that supports the multipath mitigation.

In Section 4, we present an application of the derived algorithms. We show an example of the presented

hypothesis iteration and show the results obtained when processing data collected during a government-sponsored live spoofing event with a DPA. The last section summarizes the paper's contributions and draws conclusions for future work.

2 | THEORETICAL DERIVATIONS

2.1 | Detection as a hypothesis test

Using noisy measurements to decide about the presence of a spoofing attack is a decision under uncertainty. Therefore, we cast the decision problem as a statistical hypothesis test. 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. Therefore, we follow the Neyman-Pearson paradigm that is independent of prior probabilities (Van Trees, 2001).

The hypothesis test is the solution to an optimization problem. We minimize the probability of missed detections P_{MD} while satisfying constraints on the false alert probability or statistical significance level $P_{FA_{max}}$ in Equation (1). An "alert" or "alarm" is equal to the rejection of H_0 .

$$\begin{aligned} \min_{\gamma} \quad & P(\log \Lambda(y) \geq \gamma \mid H_1) \\ \text{s.t.} \quad & P(\log \Lambda(y) < \gamma \mid H_0) \leq P_{FA_{max}}, \end{aligned} \quad (1)$$

where γ is the detection threshold and Λ is the likelihood ratio

$$\Lambda(\mathbf{y}) = \frac{p(\mathbf{y} \mid H_0)}{p(\mathbf{y} \mid H_1)}, \quad (2)$$

for the vector of measurements or evidence \mathbf{y} . An alarm is raised if $\log \Lambda(\mathbf{y}) < \gamma$.

The probabilities of Missed Detection (MD) and False Alert (FA) are represented in (1) by the terms

$$\begin{aligned} P_{MD} &= P(\log \Lambda \geq \gamma \mid H_1) \\ P_{FA} &= P(\log \Lambda < \gamma \mid H_0). \end{aligned} \quad (3)$$

Equations (1) through (3) constitute a traditional statistical hypothesis test under the Neyman-Pearson paradigm (Van Trees, 2001).

The threshold γ is found by solving the quantile function or inverse cumulative density function (cdf) of the random variable $\log \Lambda(y)$ conditioned on $y \sim H_0$ for $P_{FA_{max}}$. This can be difficult to solve analytically and might have to be done through Monte Carlo analysis offline as done

by (Psiaki et al., 2014; Rothmaier, Chen, & Lo, 2019). In the next section, we will phrase hypotheses for DoA-based approaches such that $\log \Lambda(y)$ can be approximated by a Normal distribution, leading to a straightforward solution for the detection threshold that can be calculated online by the receiver.

The requirement for the approach presented in this paper is that the measurement's behavior under nominal conditions can be exactly defined. This is not the case for dual-antenna setups used in (Borio & Gioia, 2016; Psiaki et al., 2014) if the antenna's attitude is unknown. No procedure for an online computation of a detection threshold that guarantees a maximum P_{FA} has been presented so far or is known to the authors for these approaches. We present results applying this paper's LRT to the approach in (Borio & Gioia, 2016) for known attitude in Subsection 2.6, creating a spoofing detection honoring $P_{FA_{max}}$ for a dual-antenna setup.

2.2 | Gaussian hypothesis formulation with dimensionality reduction

DoA-based spoofing detection is often formulated as a Generalized Likelihood Ratio Test (GLRT) (Van Trees, 2001), essentially a two-step process (Borio & Gioia, 2016; Konovaltsev et al., 2013; Meurer et al., 2012; Psiaki et al., 2014). First, Maximum Likelihood Estimates (MLEs) of the antenna's attitude and the spoofer's direction are estimated by aligning the measured DoAs with the expected directions and averaging over the measured DoAs respectively. In a second step, the conditional probabilities used in Equation (2) are calculated. Using the attitude estimate, the DoAs measured in an antenna fixed coordinate frame are rotated into the global frame where they can be compared to the expected directions to derive $p(\mathbf{y} \mid H_0)$. Comparing the measured DoAs to the MLE of the spoofer's direction results in $p(\mathbf{y} \mid H_1)$.

In this subsection, we will phrase augmented measurement equations independent of the nuisance parameter's attitude and spoofer direction. This results in a Normally distributed decision variable $\log \Lambda$, allowing for a fast, online computation of the decision threshold while guaranteeing a maximum false alert probability. A similar guarantee is possible with the GLRT formulated in (Konovaltsev et al., 2013; Meurer et al., 2012) but generally resulting in a higher missed detection probability as we explore by simulation in Subsection 2.5.

In phrasing augmented measurements independent of nuisance parameters, we adopt an idea followed in several papers such as Magiera and Katulski (2015) and Borio and Gioia (2016). DoAs to N satellites in view generally deliver $2N$ measurements (e.g. N azimuth, N elevation).

The unknown antenna's attitude contains three degrees of freedom, reducing the number of equations available for spoofing detection to $2N - 3$. We therefore phrase $2N - 3$ adjusted measurement equations of great circle arcs between the satellite DoAs as they are independent of the antenna's orientation (Greenwood, 1987).

The noise on each individual DoA measurement is assumed to be a rotation with zero-mean Normally distributed magnitude in an arbitrary direction as it is characterized, for example, in (Appel et al., 2015; Konovaltsev et al., 2013; Meurer et al., 2016; Meurer et al., 2012). Following the elevation-dependent over-bounding Gaussian error model characterized in (Meurer et al., 2016), we can describe the effect of noise on the measurements as a quaternion multiplication,

$$p_y = p_v \otimes p_t \quad (4)$$

$$p_v = \left[\cos \frac{\Delta_n}{2}, 0, \sin \frac{\Delta_n}{2} \sin \alpha, \sin \frac{\Delta_n}{2} \cos \alpha \right] \quad (5)$$

$$\alpha \sim U(0, 2\pi) \quad (6)$$

$$\Delta_n \sim N(0, \sigma_n^2), \quad (7)$$

where the true DoA in antenna coordinates is given by the quaternion p_t , the measured DoA by p_y and the noise by p_v . The \otimes operator denotes a quaternion multiplication. The uniformly distributed angle α defines the direction of the DoA error; Δ_n is the magnitude of the spatial angle between measured and true DoA. The noise variance σ_n^2 is characterized through the mean squared error as in (Meurer et al., 2016).

The distribution of the great circle arcs between DoAs distributed Normally as in Equations (4) through (7) is nontrivial to describe. 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 following error model over-bounds the measurement error. We start by defining a zero-mean Normal distribution with variance

$$\sigma_{ij}^2 = \sigma_i^2 + \sigma_j^2, \quad (8)$$

of an arc between DoAs i and j . Great circle arcs are further correlated. Let one great circle arc be between DoAs i and j and a second arc between DoAs j and k . We model their correlation coefficient ρ_{ijk} as

$$\rho_{ijk} = w_{ij} w_{jk} \cos \zeta_{ijk}, \quad (9)$$

where ζ_{ijk} is the spherical angle between the two arcs at DoA j under nominal conditions and w_{ij} , w_{jk} are weight vectors. The spherical angle is given by the spherical law of cosines

$$\cos \zeta_{ijk} = \frac{\cos \delta_{ik} - \cos \delta_{ij} \cos \delta_{jk}}{\sin \delta_{ij} \sin \delta_{jk}}, \quad (10)$$

with the great circle arc between DoAs i and j denoted δ_{ij} . The weight vectors are reducing the correlation coefficient for small arcs. Due to the nonnegativity of great circle arcs, the Gaussian characterization of Equation (8) is a poor approximation for small arcs and the correlation is no longer valid. The weights are calculated by

$$w_{ij} = 1 - \exp\left(-\frac{\delta_{ij}^2}{2\sigma_{ij}^2}\right). \quad (11)$$

We formalize the augmented measurement vector \bar{y} using this approximate, over-bounding error model in Equation (12). Under H_0 we expect the adjusted measurements to represent $2N - 3$ great circle arcs $\bar{\phi}$ calculated from the true azimuth and elevation values. Under H_1 the expected arcs are zero. The adjusted measurement noise $\bar{\epsilon}$ is zero-mean Gaussian with the adjusted covariance matrix \bar{R} .

$$\begin{aligned} H_0 : \bar{y} &= \bar{\phi} + \bar{\epsilon} \\ H_1 : \bar{y} &= \bar{\epsilon} \end{aligned} \quad \text{with } \bar{\epsilon} \sim N(0, \bar{R}). \quad (12)$$

Let the n th element of \bar{y} be the arc between DoAs i and j . The entry at index (n, m) of \bar{R} is then given by

$$\bar{R}_{(n,m)} = \begin{cases} \sigma_{ij}^2 & \text{if } n = m \\ \rho_{ijk} \sigma_j^2 & \text{else if } \bar{\phi}_{(m)} = \delta_{jk} \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

where $\bar{\phi}_{(m)} = \delta_{jk}$ indicates that the m th element of $\bar{\phi}$ is an arc between DoAs j and k . Matlab code to generate \bar{R} and examples comparing the test to the approach presented in (Appel et al., 2019; Appel et al., 2015) can be found at <https://github.com/stanford-gps-lab/spoofing-detection.git>.

Based on the measurement equations (12), the conditional probabilities $p(\bar{y} | H_0)$ and $p(\bar{y} | H_1)$ are evaluations of multivariate Normal distributions. Under H_0 with mean $\mu_0 = \bar{\phi}$, under H_1 with zero mean $\mu_1 = 0$ and in either case with covariance matrix \bar{R} , Equation (2) is then a special case of the general Gaussian problem (Van Trees, 2001). $\log \Lambda(\bar{y})$ develops into the well-known result of

Equation (14).

$$\begin{aligned} \log \Lambda(\bar{\mathbf{y}}) &= \log \frac{p(\bar{\mathbf{y}} | H_0)}{p(\bar{\mathbf{y}} | H_1)} \\ &= -\frac{1}{2} \left((\bar{\mathbf{y}} - \mu_0)^T \bar{\mathbf{R}}^{-1} (\bar{\mathbf{y}} - \mu_0) \right. \\ &\quad \left. - (\bar{\mathbf{y}} - \mu_1)^T \bar{\mathbf{R}}^{-1} (\bar{\mathbf{y}} - \mu_1) \right) \\ \log \Lambda(\bar{\mathbf{y}}) &= (\mu_0^T - \mu_1^T) \bar{\mathbf{R}}^{-1} \bar{\mathbf{y}} - \frac{1}{2} (\mu_0^T \bar{\mathbf{R}}^{-1} \mu_0 - \mu_1^T \bar{\mathbf{R}}^{-1} \mu_1). \end{aligned} \quad (14)$$

Substituting in $\mu_0 = \bar{\phi}$ and $\mu_1 = 0$ from Equation (12) results in:

$$\log \Lambda(\bar{\mathbf{y}}) = \bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\mathbf{y}} - \frac{1}{2} \bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi}. \quad (15)$$

Using the distribution of $\bar{\mathbf{y}}$ under either hypothesis in Equation (12) we can finally derive the distribution of $\log \Lambda(\bar{\mathbf{y}})$ under nominal and spoofed conditions:

$$\log \Lambda(\bar{\mathbf{y}}) | H_0 \sim N \left(\frac{1}{2} \bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi}, \bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi} \right) \quad (16)$$

$$\log \Lambda(\bar{\mathbf{y}}) | H_1 \sim N \left(-\frac{1}{2} \bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi}, \bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi} \right). \quad (17)$$

The optimization problem formulated in Equation (1) is now easily solved using the result in Equation (16). The detection threshold γ is given by Equation (18), where Φ^{-1} is the quantile function or inverse cdf of the Standard Normal distribution.

$$\gamma = \frac{1}{2} \bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi} + \Phi^{-1} (P_{FA_{max}}) \sqrt{\bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi}}. \quad (18)$$

We note the low computational complexity of the involved calculations, specifically of the decision threshold in Equation (18). The size of $\bar{\mathbf{R}}$ and $\bar{\phi}$ is naturally bounded by the number of satellites in view for a given constellation. $\Phi^{-1}(P_{FA_{max}})$ could be precomputed offline for a chosen false alert probability, or alternatively be implemented in the form of a lookup table for selected values of $P_{FA_{max}}$.

We can calculate the missed detection probability for a spoofing attack where a set of malicious satellite signals are radiated from the same direction for a given satellite geometry, measurement accuracy and maximum false alert probability using Equation (19).

$$\begin{aligned} P_{MD} &= 1 - \Phi \left(\frac{\gamma + \frac{1}{2} \bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi}}{\sqrt{\bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi}}} \right) \\ &= 1 - \Phi \left(\sqrt{\bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi}} + \Phi^{-1} (P_{FA_{max}}) \right). \end{aligned} \quad (19)$$

We note that we are using an approximate, overbounding error model. Any measurement including spoofed measurements $\bar{\mathbf{y}}$ will appear to match H_0 closer than if the error model were perfect. The result of Equation (19) will therefore be an optimistic estimate of the true P_{MD} .

Nevertheless, the results from Equations (18) and (19) underline the direct dependency of DoA-based detection techniques on satellite geometry and measurement accuracy. A larger value of the Mahalanobis distance $\bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi}$ represents a more powerful test. We graphically explore this relationship in Subsection 2.4.

2.3 | Selection of $2N - 3$ arcs

We have now set up augmented measurement equations to phrase hypotheses independent of the antenna's attitude and the spoofer's direction. But we have left some ambiguity. $N(N-1)/2$ great circle arcs can be spanned between N DoAs, but $\bar{\mathbf{y}}$ only comprises $2N - 3$ independent arcs such that $\bar{\mathbf{R}}$ remains full rank. If $N \geq 4$ there is more than one possible set of independent arcs that can be selected. The set of arcs S that minimizes P_{MD} is the solution to an optimization problem. Leveraging Equation (19) we can cast it as follows

$$\begin{aligned} \max_s \quad & \bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi} \\ \text{s.t.} \quad & \det \bar{\mathbf{R}} \neq 0. \end{aligned} \quad (20)$$

Equation (20) is a maximization over solutions to a Boolean satisfiability problem, which is in general NP-hard to solve (Cook, 1971).

For the algorithm to be executed online by a receiver, we therefore resort to random sampling. A sample is accepted if it is feasible in terms of the constraint and leads to a covariance matrix with a conditioning number below $10^{N/3}$. A well-conditioned covariance matrix indicates a diverse set of arcs resulting both in a numerically stable inversion and a powerful test. Around one-third of samples meet these requirements. A selection of arcs is therefore easily found. All results shown in the remainder of this paper were produced following this approach.

2.4 | Limit case of two satellites

To illustrate the performance of the detection, we show an analysis of the most challenging scenario for a detection method based on DoAs. It is the case when only two satellite signals are broadcasted from the same direction. In this case the augmented measurement vector is the scalar length of the great circle arc between the two satellites. The covariance matrix is modeled by $\bar{\mathbf{R}} = 2\sigma^2$ for the

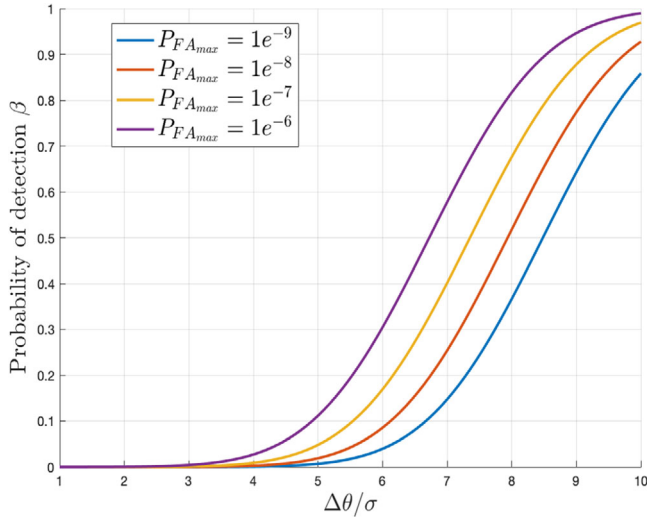


FIGURE 2 Spoofing detection probability using two satellites as a function of maximum false alert probability, satellite separation and measurement accuracy [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

measurement standard deviation σ . We show the probability of detection $\beta = 1 - P_{MD}$ calculated using Equation (19) in Figure 2. We plot it against the great circle arc $\bar{\phi} = \Delta\theta$ between the two satellites normalized by the measurement standard deviation σ for different maximum false alert probabilities. Figure 2 shows that to be able to detect a spoofer that radiates only two satellite signals from the same direction reliably, the satellite spacing should be around seven-to-ten times the measurement standard deviation.

2.5 | Comparison to a GLRT on $2N$ measurements

When using the optimal set of great circle arcs as defined by Equation (20), the simple vs. simple LRT with detection threshold given by Equation (18) is the most powerful test independent of the nuisance parameters; it presents a UMPI test (Lehmann & Romano, 2005). However, it is not strictly more powerful than a GLRT on $2N$ measurements as presented in (Konovaltsev et al., 2014; Konovaltsev et al., 2013; Meurer et al., 2016; Konovaltsev et al., 2012). We further accept a reduction in test power when choosing a suboptimal set of great circle arcs as described in the previous subsection and by employing the conservative error model described in Subsection 2.2. To validate and justify this paper's approach, we examine its missed detection probability through simulations and compare it against the approach in the cited literature.

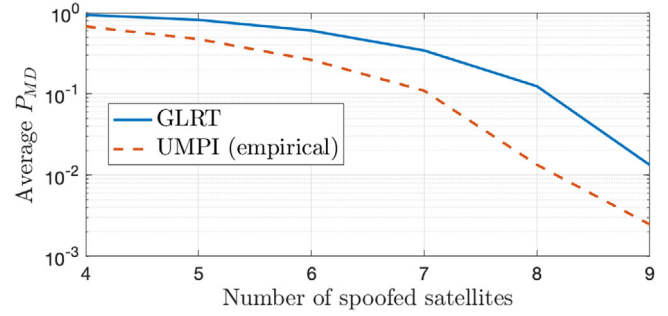


FIGURE 3 Average probability of missed detection among all subsets of n out of 9 satellites for the GLRT and UMPI test. The results for the GLRT are calculated by Equation (21); the results from the UMPI are based on Monte Carlo simulations [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

In Figure 3 we display average P_{MD} values for varying numbers of spoofed satellites. All results are based on the same constellation of nine satellites with azimuth and elevation values distributed randomly about the sky. For each number of satellites between four and nine, we compute the P_{MD} for every possible subset of satellites. The detection thresholds are set to satisfy $P_{FA_{max}} = 10^{-7}$ per measurement epoch. Testing all subsets of four to nine spoofed satellites requires a total of 382 tests. Each individual test threshold is therefore set to satisfy $P_{FA} < 10^{-7} / 382$. The DoA standard deviation is set to $\sigma_i = 12 \text{ deg}$.

The satellites not in the subset are considered neither in the test nor in the attitude estimation of the GLRT. This is meant to reflect a scenario where a spoofer is broadcasting only a subset of satellite signals from one direction but possibly broadcasts the other signals from other directions. The P_{MD} values in Figure 3 are averaged over all subsets for each number of satellites in the spoofed subset. All P_{MD} values of this paper's approach are computed through 10^6 Monte Carlo simulations, as the analytic expression in Equation (19) is optimistic. The P_{MD} of the GLRT is calculated by

$$P_{MD} = P_{2N-3, \lambda} \left(C_{2N-3}^{-1} \left(1 - 10^{-7} / 382 \right) \right), \quad (21)$$

where C_k^{-1} is the inverse chi-squared cdf with k degrees of freedom. $P_{k, \lambda}$ is the cdf of the noncentral chi-squared distribution with k degrees of freedom and noncentrality parameter λ given by (Konovaltsev et al., 2014).

$$\lambda = \sum_{n=1}^N \left(\frac{\psi_n}{\sigma_n} \right)^2 \quad (22)$$

where ψ_n is the spatial angle between the n th predicted ephemeris-based DoA and spoofed DoA as in

(Konovaltsev et al., 2014; Konovaltsev et al., 2013; Meurer et al., 2016; Meurer et al., 2012).

As we can see in Figure 3, in the simulated scenario the UMPI test presented in this paper outperforms the GLRT from the cited literature with a more than 50% lower missed detection rate. This result is just one qualitative comparison of the approaches under the described attack scenario. The underlying code is accessible at <https://github.com/stanford-gps-lab/spoofing-detection.git> for the interested reader and to facilitate further comparisons.

2.6 | Application to a dual-antenna setup with known attitude

Setups of two closely spaced antennas are capable of deriving a spatial metric by measuring the difference in carrier phase (Borio & Gioia, 2016; Psiaki et al., 2014). The following brief derivation leverages the approximations justified in (Borio & Gioia, 2016) and equally considers the integer ambiguities as random variables. Specifically, the spatial angle between a line connecting the two antenna phase centers and the direction of the signal from the i th satellite is characterized by

$$\Delta \varphi_i = \frac{D}{\lambda} \cos \alpha_i + \Delta N_i + \frac{c}{\lambda} (dT^2 - dT^1) + \frac{1}{\lambda} \Delta \eta_i, \quad (23)$$

where $\Delta \varphi_i$ is the single difference in carrier phase in units of cycles, D is the distance between the antenna phase centers, λ is the signal wave length, α_i is the spatial angle, ΔN_i is the cycle ambiguity as an integer, c is the speed of light, dT^j is the clock error of the j th receiver. The noise on the i th measurement $\Delta \eta_i$ is the i th element of the noise vector $\Delta \eta$.

$$\Delta \eta \sim N(0, R) \text{ with } R = \text{diag}(\sigma_1^2, \dots, \sigma_n^2) \quad (24)$$

A procedure to characterize the σ_i 's is given in (Borio & Gioia, 2016).

Under nominal conditions, the spatial angle is different for every satellite. Spoofed satellite signals coming from the same direction have the same spatial angle. No threshold computation guaranteeing a $P_{FA_{max}}$ has been presented in the literature for this setup, nor is one possible with this paper's approach, due to the nonlinearity of the cosine function.

If the antenna's attitude is known, however, this paper's approach applies. Single differences of N satellites deliver N independent measurements. The unknown spoofer's direction forms one nuisance parameter. Therefore, we phrase $N - 1$ double difference equations as the differences between carrier phase single difference mea-

surements. The vector of $N - 1$ double differences is given by

$$\Delta^2 \varphi = A \Delta \varphi = \frac{D}{\lambda} A \cos \hat{\alpha} + A \Delta N + \frac{1}{\lambda} A \Delta \eta, \quad (25)$$

where $\hat{\alpha}$ represents the vector of expected spatial angles given the known or estimated attitude. The matrix A generates differences between the i th and j th measurement. Without loss of generality we can assume the measurements in $\Delta \varphi$ to be sorted with increasing $\hat{\alpha}_i$. A is then given by

$$A = \begin{bmatrix} -1 & 1 & \cdots & 0 & 0 \\ 0 & -1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & \cdots & -1 & 1 \end{bmatrix} \in R^{N-1 \times N}. \quad (26)$$

We can now define the behavior under nominal and spoofed conditions, where spoofed conditions assume all signals being broadcasted from the same direction.

$$\begin{aligned} H_0 : f(\Delta^2 \varphi) &= \frac{D}{\lambda} A \cos \hat{\alpha} + \frac{1}{\lambda} \Delta^2 \eta \quad \text{with } \Delta^2 \eta \sim N(0, A R A^T), \\ H_1 : f(\Delta^2 \varphi) &= \frac{1}{\lambda} \Delta^2 \eta \end{aligned} \quad (27)$$

where the function $f(x)$ is defined as

$$f(x) = x - \text{round}(x). \quad (28)$$

The attitude information is only known with finite precision. An error in attitude ε_a causes an error in carrier phase single difference that is given in the unit of length by

$$\varepsilon_{SD} = D(\cos(\alpha + \varepsilon_a) - \cos \alpha). \quad (29)$$

For small attitude errors we can approximate this error by a first order Taylor expansion and further simplify the expression leveraging the small angle assumption.

$$\begin{aligned} \varepsilon_{SD} &= -D \sin(\alpha + \varepsilon_a) \varepsilon_a \\ &= -D(\sin \alpha + \cos \alpha \varepsilon_a) \varepsilon_a \\ &\approx -D \sin \alpha \varepsilon_a. \end{aligned} \quad (30)$$

Thanks to this approximately linear relationship, we can now model uncertainty in attitude by an elevated uncertainty in carrier phase single differences. We define and work with the updated variance of the i th single difference measurement $\tilde{\sigma}_i^2$.

$$\tilde{\sigma}_i^2 = \sigma_i^2 + (D \sin \alpha_i)^2 \sigma_a^2. \quad (31)$$

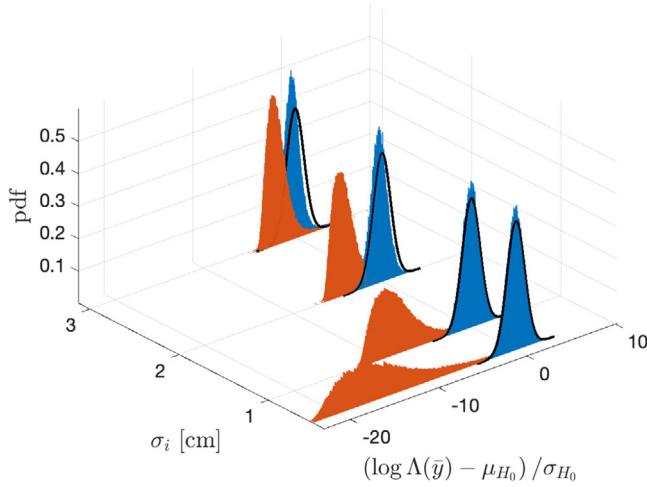


FIGURE 4 Pdfs estimated through Monte Carlo simulations and theoretical measurement model for different measurement standard deviations σ [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

To examine this algorithm's performance and validate the error model, we run Monte Carlo simulations for a scenario of $N = 6$ satellites with randomized geometries. We set $\sigma_a = 2^\circ$, a realistic root mean squared error for a small low-cost attitude determination system (Hyyti & Visala, 2015). Similar to (Borio & Gioia, 2016), we consider values of σ_i between 0.5 cm and 3 cm. In Figure 4, we compare the empirical pdfs of $\log \Lambda$ for nominal (blue) and spoofed (red) measurements, all normalized by the mean and standard deviation of $\log \Lambda$ under nominal conditions as given by (16). Black lines are Standard Normal pdfs, representing the measurement model. We can see the measurements adhering to the Gaussian model under nominal conditions. The randomized geometry flattens the distribution under spoofed conditions.

In Figure 5 we show the Receiver Operating Characteristics (ROCs) (Van Trees, 2001) for various values of σ_i to facilitate a comparison with (Borio & Gioia, 2016), again with $\sigma_a = 2^\circ$. We obtain high values of the detection probability P_D already for small values of $P_{FA_{max}}$. Adding attitude information results in higher detection performance than in (Borio & Gioia, 2016) and more importantly allows for the detection threshold being set to guarantee a constraint on false alerts.

3 | HYPOTHESIS ITERATION

Except for the analysis in Section 2.5, the derivations so far have only considered the simple “all satellites spoofed from the same transmitter” or “all satellites nominal” hypothesis. We have shown in (Rothmaier et al., 2019)

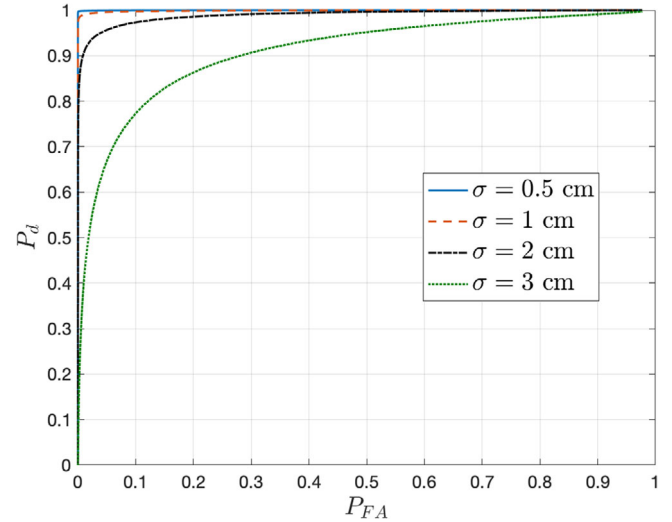


FIGURE 5 ROC curves for attitude uncertainty $\sigma_a = 2$ deg for different carrier phase measurement standard deviations σ [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

that conditions are rarely this clean under real-world conditions. In the following, we present strategies to deal with violations to either simple hypothesis that represent updated versions of our approach presented in (Rothmaier et al., 2019). We suggest applying all of the following procedures to each GNSS constellation separately, as a separate attack has to be launched for each constellation.

The presented techniques further come without performance guarantees but rather are computationally cheap heuristics that have worked well in our experience.

3.1 | Nominal conditions

Under nominal conditions, DoA measurements are affected by multipath as reported, for example, in (Konovaltsev et al., 2013) for an antenna array or in (Rothmaier et al., 2019) and (Egea-Roca et al., 2018) for a Dual Polarization Antenna (DPA), thereby violating the assumption of zero-mean Normally distributed measurement errors.

To find a tradeoff between mitigating the effect of multipath on false alerts without dramatically reducing the detection capability of the test, we leverage another integrity algorithm already in place: Receiver Autonomous Integrity Monitoring (RAIM). In its aviation implementation, it detects faults on single GPS satellites, while Advanced RAIM (ARAIM) is designed to capture faults on multiple satellites and constellation-wide faults by using satellites from many constellations (Blanch et al., 2014). We can safely exclude a single GPS satellite from the hypothesis test, since protection against wrong information from a single satellite is provided by RAIM. At each epoch we

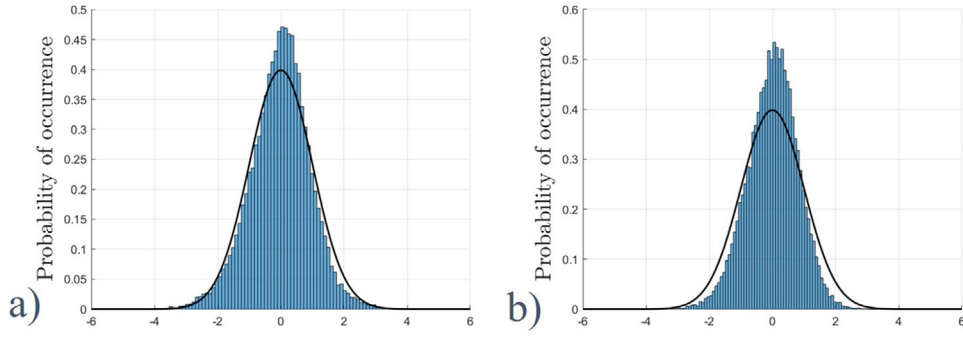


FIGURE 6 Normalized azimuth measurement error in flight from a Dual Polarization Antenna (DPA) mounted on a C12 aircraft. In a) for all measurements, in b) for the measurements after the multipath mitigation. The black line is a Standard Normal distribution and represents the measurement model [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

exclude the satellite from the computation that has the most negative contribution to $\log \Lambda(\bar{\mathbf{y}})$ when normalized using its distribution under nominal conditions given by Equation (16) (and therefore the largest contribution to an alert). Depending on how $\bar{\phi}$ is calculated, this can be done using an analytical expression of the gradient of the normalized $\log \Lambda(\bar{\mathbf{y}})$ w.r.t. each satellite. Given the simplicity of the involved calculations and small domain of the problem, a global search over all N versions of the normalized $\log \Lambda(\bar{\mathbf{y}})$, each with one satellite removed, can be done without problem. The index of the satellite to be excluded i^* is formally determined using Equation (32), where the subscript i denotes variables without the excluded i th satellite.

$$i^* = \arg \max_i \left(\frac{\log \Lambda(\bar{\mathbf{y}}_i) - \frac{1}{2} \bar{\phi}_i^T \bar{\mathbf{R}}_i^{-1} \bar{\phi}_i}{\sqrt{\bar{\phi}_i^T \bar{\mathbf{R}}_i^{-1} \bar{\phi}_i}} \right) \\ = \arg \max_i \left(\frac{\bar{\phi}_i^T \bar{\mathbf{R}}_i^{-1} \bar{\mathbf{y}}_i}{\sqrt{\bar{\phi}_i^T \bar{\mathbf{R}}_i^{-1} \bar{\phi}_i}} - \sqrt{\bar{\phi}_i^T \bar{\mathbf{R}}_i^{-1} \bar{\phi}_i} \right). \quad (32)$$

A new set of arcs will have to be chosen for most of the N subsets as outlined in Subsection 2.3.

Removing a single satellite is not sufficient to compensate for the effect of multipath in a dense urban environment. While this technique can still be used under these conditions, the Gaussian error model likely will not over-bound the measurement errors. This results in a false alert probability higher than what is guaranteed. While other mitigations may be possible under such conditions, robust performance has yet to be demonstrated for any DoA-based technique. This paper focuses on first solving the more benign situation with few faults.

As an application example, in Figure 6, we show the normalized error of azimuthal DoA measurements from a DPA mounted on a C12 aircraft in flight, both before and after the multipath mitigation step. The flight profile included various climbs, descends and turns with up to 60 deg bank. An extensive description of the data collection campaign is given by (Fulton et al., 2020). Assuming coordinated turns, pitch and bank are roughly estimated from the GNSS velocity vector. Based on 5 hours of flight data, the measurement standard deviation is inflated as a linear function of bank angle θ . (Rothmaier et al., 2019) describes in detail how to estimate the measurement standard deviation σ for the DPA. After inflation due to the aircraft's bank angle, it is given by

$$\tilde{\sigma} = \sigma \left(1 + \frac{\theta}{60^\circ} \right). \quad (33)$$

The black line shows a Standard Normal probability density function (pdf) for comparison, representing the measurement model including inflated measurement uncertainty. Before the mitigation step, the measured error distribution has tails stronger than a Standard Normal distribution; after the mitigation the Standard Normal distribution over-bounds the measurement error. This guarantees a conservative alert behavior that satisfies the false alert probability constraint and validates the multipath mitigation strategy.

3.2 | Spoofed conditions

A more dramatic simplification is inherent to the “all satellites spoofed from the same transmitter” hypothesis. Violations of the assumption that all signals will come from the same direction under spoofed conditions can easily be imagined. Either the receiver has not locked onto all

signals emitted by the spoofer and is still tracking some authentic satellite signals, or the attacker is using multiple antennas to transmit signals from multiple sources. We have shown in (Rothmaier et al., 2019) that the first is often the case, reducing the detection capability if it is not accounted for. And while the latter attack scenario is difficult to mount, it is possible – especially for the case of a cooperative victim (Psiaki & Humphreys, 2016).

Few DoA-based approaches specifically deal with this simplification. Reference (Psiaki et al., 2014) briefly suggests an M-ary hypothesis test, thereby increasing the complexity of calculating a robust decision threshold. (Meurer et al., 2012) already presents an interesting iterative approach but is constrained by computational limitations due to a necessary attitude computation for each possible subset and the large number of subsets. Reference (Esswein & Psiaki, 2019) examines the issue more extensively and suggests an exhaustive search over all $(1 + S)^N$ possible combinations of nominal and spoofed subsets for N satellites and S possible spoofer antennas. This comes with two drawbacks. One is the obvious computational load associated with the large number of possible combinations that are being considered, especially since a nonconvex optimization problem is solved for each combination. The second drawback is the more subtle aspect that to guarantee a certain false alert probability per measurement epoch, the false alert probability used to compute the threshold as in Equation (18) has to be adjusted by a factor equal to the number of combinations considered at that epoch. We have done so in our results presented in Subsection 2.5. Testing for several hundred hypotheses in parallel may not be a computational burden for future receivers. To nevertheless offer a computationally cheaper alternative, we show an updated version of the greedy iterative algorithm to find the largest subsets that justify raising an alarm introduced in (Rothmaier, Chen, & Lo, 2019). The algorithm starts by testing all N satellites of a constellation in view. If no alarm is raised, all subsets of $N - 1$ satellites are examined and the one resulting in the lowest normalized $\log \Lambda(\bar{\mathbf{y}})$ is selected. If it does not cause an alarm, all $N - 2$ subsets reachable by removing one more satellite from the selected set are examined. Again, the subset resulting in the lowest normalized $\log \Lambda(\bar{\mathbf{y}})$ is selected and tested. Satellites are removed in this greedy manner until an alarm is raised or no more satellites remain.

To ensure the validity of the measurement model, we further apply the multipath mitigation step of (32) before testing $\log \Lambda(\bar{\mathbf{y}}_{i^*}) < \gamma_{i^*}$. If the test is negative, the i th satellite is added back to the consideration before the next subset reduction.

While this algorithm might seem cumbersome, it only considers a limited number of satellite subsets for a minimal computational load while ensuring the validity of

the measurement model at every step. It follows the following pseudocode:

1. Start with set \mathbf{S} of all satellites in view
2. While $|\mathbf{S}| \geq$ minimum number of satellites
 - a. Remove i th satellite using Equation (32)
to ensure error bound model
 - b. Calculate $\log \Lambda(\bar{\mathbf{y}})$, γ
 - c. If $\log \Lambda(\bar{\mathbf{y}}) < \gamma$ break; else
stop if current set \mathbf{S} causes alarm
 - d. Remove j th satellite using Equation (34)
next smaller subset, this reduces $|\mathbf{S}|$ by 1

$2N - 3$ adjusted measurement equations are used to compute $\log \Lambda(\bar{\mathbf{y}})$ in Step 2.b, requiring a minimum of $N = 2$ satellites. One is removed in the multipath consideration Step 2.a, constraining the theoretical minimum number of satellites for Step 2 to $N_{\min} = 3$. A higher value can be used to reduce, for example, the computational load.

The main step is the removal of a satellite in Step 2.d. Analogous but opposite to the satellite selection due to multipath, we now remove the satellite whose removal leads to the smallest normalized $\log \Lambda(\bar{\mathbf{y}})$. The index of the satellite to be removed j^* is formally determined using Equation (34), where the subscript j denotes a variable after the exclusion of the j th satellite.

$$j^* = \arg \min_j \left(\frac{\bar{\phi}_j^T \bar{\mathbf{R}}_j^{-1} \bar{\mathbf{y}}_j}{\sqrt{\bar{\phi}_j^T \bar{\mathbf{R}}_j^{-1} \bar{\phi}_j}} - \sqrt{\bar{\phi}_j^T \bar{\mathbf{R}}_j^{-1} \bar{\phi}_j} \right). \quad (34)$$

We note the small number of subsets that are overall considered. For subsets of K spoofed satellites among a total of N satellites, $1 + \sum_{k=1}^{N-K} (N - k + 1) = 1 + \frac{1}{2}(N^2 + N - (K^2 + K))$ subsets are considered. In the above example of $N = 10$ satellites and the extreme case of $K = 2$ spoofed satellites, 53 subsets are considered, while the exhaustive global search of (Esswein & Psiaki, 2019) evaluated 1,024 combinations.

This greedy minimization algorithm does not necessarily find the subset with the globally smallest normalized $\log \Lambda(\bar{\mathbf{y}})$. In return, it comes with significantly reduced computational complexity. Results on test data from a government-sponsored live spoofing event that we present in the next section show strong results of this approach, despite the possibility of working with subsets that represent local instead of global minima. The best choice of approach will depend on the receiver's capabilities and application and is up to the designer.

We summarize the algorithm presented in the last two subsections in a block diagram in Figure 7, starting from

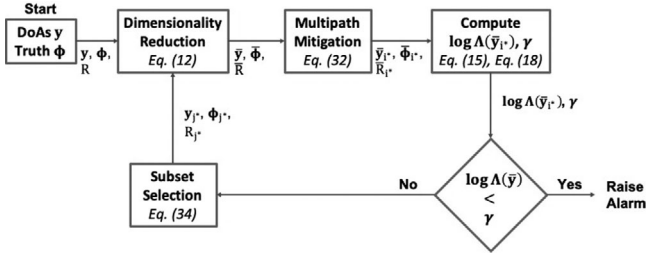


FIGURE 7 Block diagram summarizing the LRT with dimensionality reduction and hypothesis iteration

the measured DoAs \mathbf{y} with measurement uncertainty \mathbf{R} and the ephemeris-based directions ϕ . We reference the key equations used in each step.

4 | APPLICATION EXAMPLE USING AZIMUTH-ONLY MEASUREMENTS

In this section, we apply the algorithmic derivations in the previous sections to data collected with a Dual Polarization Antenna (DPA) during a government-sponsored live spoofing event. It is the antenna architecture that was used to collect the flight data shown in Figure 6 and that has been presented in (Chen et al., 2018; Chen et al. 2017; McMilin, 2016). Specifically, we define the measurement model for azimuth-only DoAs, show an example of the presented hypothesis iteration and give statistics on the algorithm's overall performance.

4.1 | Problem formulation using azimuth-only measurements

The DPA architecture developed at Stanford delivers azimuthal DoA measurements that can be used for spoofing detection (Chen et al., 2018; Chen et al., 2017; Lo et al., 2020; McMilin, 2016; Rothmaier, Chen, Lo, & Powell, 2019). At the price of increased measurement noise, the DPA delivers DoAs from a single element antenna. The DoA computation is based on the phase difference between Right Hand and Left Hand Circular Polarized (RHCP and LHCP, respectively) signals. (Lo et al., 2020) presents the derivation of azimuth measurements from the phase difference in more detail. GNSS signals are generally RHCP signals, but a significant LHCP signal component has been observed by several groups (Egea-Roca et al., 2018; Esswein & Psiaki, 2019). Due to the azimuth determination procedure of the DPA, the measurements come with a 180 deg ambiguity.

Azimuth measurements are affected by the antenna's attitude. For most aircraft we can assume coordinated turns and a pitch angle similar to the flight path angle. Pitch and bank can then be roughly estimated from the GNSS velocity vector, as we have done for the results shown in Figure 6. Similar considerations can be made for many applications on land vehicles. Ships in heavy seas or highly maneuverable fighter aircraft do not allow for these considerations and violate small angle assumptions on pitch and bank. The setup considered in the remainder of this paper is not applicable to these applications without additional information, for example, from an IMU.

With pitch and bank estimated or constrained we then phrase an augmented measurement vector $\bar{\mathbf{y}}$ of size $N - 1$ that contains the differences of N DoA azimuth measurements as defined in Equation (35). The expected measurements under nominal conditions are similarly the differences of azimuth directions ϕ of the satellites.

$$\begin{aligned}\bar{y}_1 &= y_2 - y_1 \\ \bar{y}_2 &= y_3 - y_2 \\ &\vdots \\ \bar{y}_{N-1} &= y_N - y_{N-1}.\end{aligned}\tag{35}$$

Without loss of generality we assume that \mathbf{y} and ϕ are sorted according to increasing true azimuths. We formulate the adjusted measurement vector after the dimensionality reduction under either hypothesis in Equation (36).

$$\begin{aligned}H_0 : \bar{\mathbf{y}} &= \mathbf{A}\bar{\phi} + \bar{\epsilon} \\ H_1 : \bar{\mathbf{y}} &= \bar{\epsilon}\end{aligned}\quad \text{with } \bar{\epsilon} \sim \mathcal{N}(0, \bar{\mathbf{R}}); \bar{\mathbf{R}} = \mathbf{A}\mathbf{R}\mathbf{A}^T,\tag{36}$$

where the matrix \mathbf{A} is defined by Equation (26) and \mathbf{R} is the covariance matrix of the original DoA measurements. Its derivation for the DPA is presented in (Rothmaier, Chen, & Lo, 2019).

Due to the 180 deg ambiguity of the DPA, the vectors $\bar{\mathbf{y}}$ and $\bar{\phi}$ generally consist of values in the range $[-\pi/2, \pi/2)$. However, for $\bar{\mathbf{y}}$ to be Normally distributed about $\bar{\phi}$, some of its values need to be corrected by a multiple of π . We illustrate this with a simple example using scalar values. Let $\bar{\phi} = 1.5$ and $\bar{\mathbf{y}} = -1.5$ with $\bar{\mathbf{R}} = \sigma^2 = 0.1$. Likely $\bar{\mathbf{y}} = -1.5 + \pi \approx 1.64$ is the correct value for the random variable, as it is less than 1σ away from its expected value $\bar{\phi}$ (the original value -1.5 is more than 9σ away). These corrections can be easily applied as long as $\sigma \ll \pi$. Without the 180 deg ambiguity the same logic would apply for the range $[-\pi, \pi)$, as it does for any other modded random variable. The interested reader is referred to the discussion in Section 4 of (Borio & Gioia, 2016) where the situation is addressed for variables in a range of $[-0.5, 0.5)$.

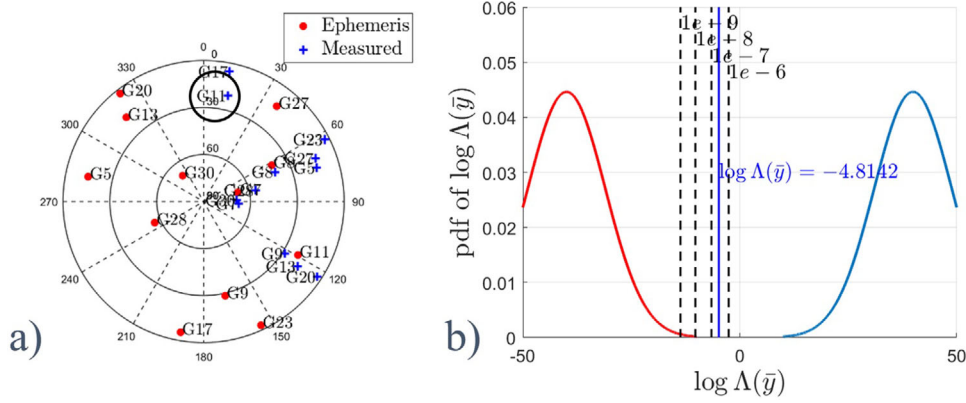


FIGURE 8 Skyplot in a) and decision variable space in b) for all GPS satellites in view. Satellite G11 is removed by the multipath mitigation step. Figure b) shows pdfs of expected distributions under H_0 (blue) and H_1 (red), thresholds depending on the maximum false alert probability as dashed vertical lines and the measured value as solid line [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

This procedure will inevitably at some point apply a correction incorrectly, as the more probable option is not always the right one. Any error introduced moves the measurement closer to its mean under nominal conditions and therefore likely leads to an increase in $\log \Lambda(\bar{\mathbf{y}})$, thereby ensuring that the constraint on a maximum false alert probability is met. To be conservative and rule out the corner case where this is not the case, we could opt to only apply the correction if it in fact leads to a higher value of $\log \Lambda(\bar{\mathbf{y}})$.

It is difficult to exactly quantify the impact of this procedure on the number of alerts under spoofed conditions, but as long as $\sigma \ll \pi$ holds, the impact is expected to be small. Examining the results with and without this correction on the data from the live spoofing conditions has shown limited impact even for considerable measurement uncertainties of $\sigma \leq \pi/6$.

Applying Equation (36) to Equations (15) and (18) results in the straightforward formulation of the log likelihood ratio in Equation (37) and detection threshold in Equation (38).

$$\log \Lambda(\bar{\mathbf{y}}) = \phi^T \mathbf{A}^T \bar{\mathbf{R}}^{-1} \bar{\mathbf{y}} - \frac{1}{2} (\phi^T \mathbf{A}^T \bar{\mathbf{R}}^{-1} \mathbf{A} \phi) \quad (37)$$

$$\gamma = \frac{1}{2} \phi^T \mathbf{A}^T \bar{\mathbf{R}}^{-1} \mathbf{A} \phi + \Phi^{-1}(P_{FA_{max}}) \sqrt{\phi^T \mathbf{A}^T \bar{\mathbf{R}}^{-1} \mathbf{A} \phi}. \quad (38)$$

4.2 | The hypothesis iteration

To illustrate the greedy hypothesis iteration algorithm described in the previous subsection, we present an example of a successful detection of a spoofed subset during the government-sponsored live spoofing event. All

azimuth measurements shown in the following plots come with standard deviations between 15 and 30 deg. Figure 8a) shows a skyplot for all GPS satellites in view. The ephemeris-based satellite positions are marked as red circles on the skyplot; positions based on the ephemeris-based elevation but estimated azimuth directions are shown as blue +. The absolute azimuth values are of no importance in the spoofing detection, as only the difference between measurements is considered. Figure 8a) shows azimuths rotated by an MLE of the antenna's heading for illustration. The MLE is calculated by solving the segmented optimization problem neatly described in Appendix A of (Borio & Gioia, 2016).

We can see some satellites detected to come from very similar azimuths (G5, G7, G8, G23, G27, G28, G30), matching the spoofed hypothesis. Others (G9, G11, G13, G17 and G20) are coming from different directions, matching the nominal hypothesis. We circle satellite G11 to indicate that has been identified for exclusion by the multipath detection algorithm presented in the previous subsection. It is ignored in the spoofing detection algorithm and in the values shown in Figure 8b). Figure 8b) shows the $\log \Lambda(\bar{\mathbf{y}})$ decision space. We plot pdfs of $\log \Lambda(\bar{\mathbf{y}})$ under nominal (blue, right-hand side) and spoofed (red, left-hand side) conditions. Vertical dashed lines represent detection thresholds for different values of $P_{FA_{max}}$. The blue solid vertical line indicates the value of $\log \Lambda(\bar{\mathbf{y}})$ given the measured azimuth values. We can see a range of aspects from these two figures:

- The distributions expected under H_0 and H_1 (blue and red curve) of $\log \Lambda(\bar{\mathbf{y}})$ are well-separated. This corresponds to a large Mahalanobis distance $\bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi}$ and a powerful detection test as we would expect from a full sky with $N = 12$ satellites in view.

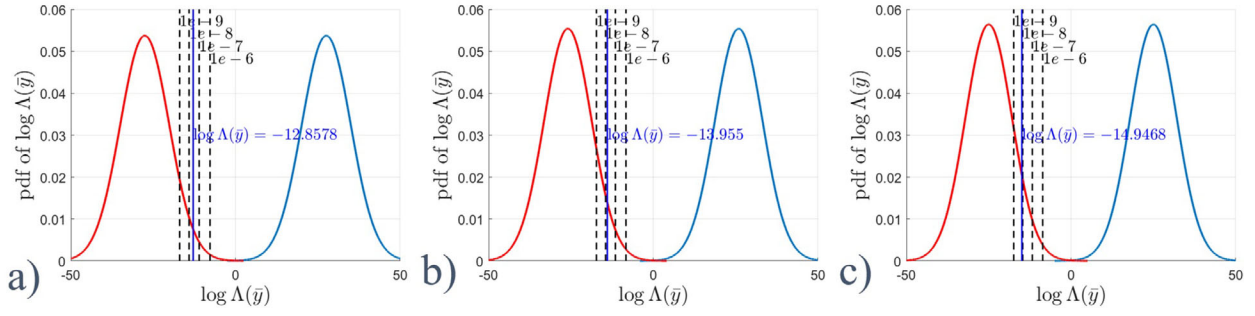


FIGURE 9 Decision spaces for consecutive removal of satellites G17, G20 and G13. Satellite G11 is removed by the multipath step after each satellite removal. The decision variable $\log \Lambda(\bar{\mathbf{y}})$ moves further left with each removal, until it finally crosses the 10^{-8} line in subplot c) [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

- The measurement \mathbf{y} and its associated $\log \Lambda(\bar{\mathbf{y}})$ match neither expected distribution well. This reflects the circumstance that \mathbf{y} is a mix of spoofed and nominal measurements.
- We are confident enough to raise an alarm if the maximum false alert probability is 10^{-6} , but not if it is 10^{-7} or lower; the blue line is to the left of the black 10^{-6} line but to the right of the 10^{-7} line.

In Figure 9 a) through c) we show the $\log \Lambda(\bar{\mathbf{y}})$ decision spaces for the next three iterations of the subset selection, removing satellites G17, then G20, and finally G13. G11 is selected in each case for exclusion in the multipath mitigation step. We can see the theoretical distributions under H_0 and H_1 move closer together with each satellite exclusion, corresponding to a smaller Mahalanobis distance $\bar{\phi}^T \bar{\mathbf{R}}^{-1} \bar{\phi}$ after the removal of a satellite. At each step, $\log \Lambda(\bar{\mathbf{y}})$ moves further left and crosses the 10^{-8} line after the third removal.

At this point an alarm is raised. A subset of satellites has been identified for which the null hypothesis can be rejected with a false alarm probability of less than 10^{-8} per measurement.

Different actions are possible when raising an alarm for a subset of satellites. The most conservative course of action is to declare the entire constellation unusable. Especially if only a small subset of satellites has been identified as spoofed, the entire procedure could be run again on the remaining satellites, to find additional spoofing sources or hopefully identify satellites as useable despite a spoofed subset.

In the latter course of action, we recommend changing the philosophy behind selecting the detection threshold. As a spoofer has already been identified, a constraint on false alert probability is no longer sensible. The goal of recursively searching for additional spoofed subsets rather has the inverse motivation: to find subsets of satellites that can be trusted despite the presence of a spoofing attack. Therefore, the threshold should be set to satisfy a con-

straint on maximum missed detection probability by solving the quartile function of the distribution of $\log \Lambda(\bar{\mathbf{y}})$ under H_1 given by Equation (17).

4.3 | Performance statistic under live spoofing conditions

We tested the algorithm developed in the past two sections on GPS and GLONASS data collected during a live spoofing event sponsored by the United States government. Data was recorded during a total of 39 episodes of spoofing, always from a single source transmitter. Each constellation was processed separately, for a total of 442 measurement epochs, of which 129 were spoofed during the 39 episodes of consecutive spoofing. During most spoofed epochs, we received a mix of genuine and spoofed signals.

To underline the necessity of the hypothesis iteration but also show results for options of different computational efforts, we show the results for three different processing methods. Method 1 only considers the “all nominal” vs. “all spoofed” hypothesis. Method 2 allows for the exclusion of one satellite from the computation to mitigate the effect of multipath as presented in Subsection 3.1. Method 3 takes full advantage of the algorithm presented in Section 3, iterating on subsets down to four satellites and mitigating multipath at each step.

The following results are for a maximum false alert probability of $P_{FA_{max}} = 10^{-7}$. Table 1 shows the numerical results for the three methods. For each method, we indicate the detected epochs and the number of episodes of spoofing during which at least one alarm was raised. We also note the number of false alarms and, more importantly, the smallest $\log \Lambda(\bar{\mathbf{y}}) - \gamma$ normalized by its expected standard deviation under nominal conditions σ_{H_0} . Epochs with $\log \Lambda(\bar{\mathbf{y}}) - \gamma < 0$ result in an alarm. How far below or above indicates the margin by which a (false) alarm was raised or not.

In Figure 10 we show histograms of the same $\log \Lambda(\bar{\mathbf{y}}) - \gamma$ values normalized by σ_{H_0} for the three approaches,

TABLE 1 Result summary using three different processing approaches on the data collected during the live spoofing event. One hundred twenty-nine epochs were spoofed during 39 episodes of consecutive spoofing

	Method 1 Binary Hypotheses	Method 2 Binary Hypotheses, Multipath Mitigation	Method 3 Hypothesis Iteration, Multipath Mitigation
Detections	29 episodes / 57 epochs	18 episodes / 30 epochs	25 episodes / 47 epochs
<i>False Alerts</i>	24	0	0
<i>Min log $\Lambda(\bar{y}) - \gamma$ while nominal</i>	$-8.33 \sigma_{H_0}$	$1.57 \sigma_{H_0}$	$1.45 \sigma_{H_0}$

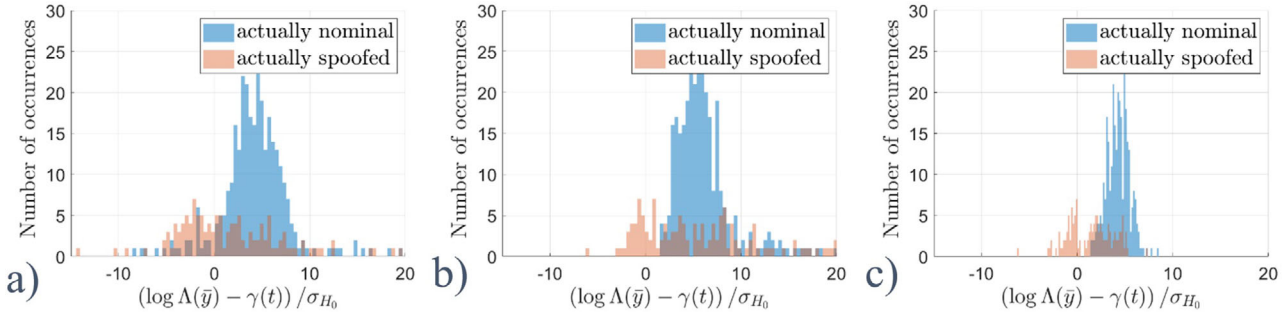


FIGURE 10 Distributions of the normalized decision variable when using Method 1 (left), Method 2 (center) and Method 3 (right). An alarm is raised for any epoch with a negative value [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

separated into nominal epochs (blue bars) and spoofed epochs (red bars). The values of γ include the inflation of $P_{FA_{max}}$ during the subset iteration algorithm described in the previous section.

The majority of measurements labelled as “spoofed” in the experiment contain a mix of spoofed and nominal satellites, creating a challenging detection scenario. The DPA comes with a simple architecture, small size and cost, but its azimuth measurements have a standard deviation of 15–30 deg with a periodicity of only 180 deg. This combination naturally results in a significant number of missed detections. Therefore, the presented results should be compared on a relative basis to evaluate the success of the presented algorithms in a very challenging detection scenario.

The first method alarms during 57 out of 129 spoofed epochs, but also raises 24 false alarms. The smallest value of $\log \Lambda(\bar{y})$ encountered during nominal conditions is $8.33 \sigma_{H_0}$ below the threshold, indicating a very strong confidence of the algorithm that spoofing is present. This poor performance is visible in the left histogram of Figure 10, where 24 nominal epochs are depicted below 0 corresponding to false alarms. Actually spoofed and actually nominal epochs are overall not very well-separated. The recorded data clearly does not match the measurement models for “all nominal” and “all spoofed” well.

The second method has fewer detections, alarming only in 30 out of 129 spoofed epochs. It raises no more false alarms; the smallest $\log \Lambda(\bar{y})$ is $1.57 \sigma_{H_0}$ above the threshold. Removing one satellite from the consideration at each epoch has reduced the number of detections noticeably

but made a dramatic difference toward avoiding false alarms. This can be observed in the middle histogram of Figure 10, where the nominal values are now comfortably above 0. However, the spoofed epochs are still not very well-separated from the nominal ones.

The third method alarms correctly 47 out of 129 epochs without raising a false alarm, a significant increase. The lowest $\log \Lambda(\bar{y})$ during nominal conditions has barely changed to $1.45 \sigma_{H_0}$ above the threshold. On the right histogram of Figure 10 the truly nominal and truly spoofed cases now seem much better separated with less extreme values. The truly nominal cases follow a distribution that roughly resembles a Gaussian distribution, as it is modelled. The subset iteration algorithm has increased the number of detections by more than 50% without noticeably impacting the guarantee on false alarms. Introducing the more realistic hypotheses has led to a better identification of both nominal and spoofed cases.

A significant number of attacks remained undetected even when using Method 3. These were challenging detection scenarios where only a small number of closely aligned satellites were spoofed. Given the large measurement uncertainty of the DPA, rejecting H_0 is not possible in these cases. This circumstance is reflected in a high expected probability of missed detection as given by Equation (19). Among all undetected epochs, the average expected P_{MD} is 79%. More than two-thirds of the undetected epochs correspond to P_{MD} values above 80%. Knowing of its low detection power, the detector can at least inform the user that it is unable to protect against spoofing.

Given the stochastic nature of the spoofing detection algorithm, the probability of at least one alert increases as more measurements are taken. An attack is more likely to be discovered the longer it persists. As we outlined in the previous section, once the attack is detected the philosophy behind setting the detection threshold changes to a constraint on maximum missed detection probability.

5 | SUMMARY AND CONCLUSIONS

This paper presents an algorithmic framework for direction of arrival (DoA)-based spoofing detection. We phrase hypothesis under dimensionality reduction to be used in the Uniformly Most Powerful Invariant (UMPI) test independent of the nuisance parameters antenna attitude and spoofer direction. The resulting detection threshold computation is highly tractable and can be done online by the receiver. We demonstrate that the algorithm outperforms previous approaches by performing simulations very similar to cited literature.

We present a hypothesis iteration algorithm that efficiently breaks the M-ary hypothesis test of detecting a spoofed subset of satellites or multiple spoofing sources down into a sequence of binary tests and mitigates the effect of weak multipath. We demonstrate flight test data of a Dual Polarization Antenna (DPA) that validates a Gaussian measurement model after the multipath mitigation.

We finally present results of the algorithm when used on noise-affected measurements taken during a government-sponsored live spoofing event. The hypothesis iteration algorithm identifies many spoofed subsets of satellites using low-quality azimuth-only DoA measurements while guaranteeing a maximum false alert probability.

Future work includes examining the algorithm's performance under more severe multipath conditions, to determine at what point the considerations around the "nominal" hypothesis from Section 3.1 are insufficient.

Most missed detections during the live spoofing event were in situations where a subset of satellites that is closely spaced in the sky was spoofed. This emphasizes a limitation of any signal-geometry-based detection approach. For a more powerful detector, sequential detection algorithms and the inclusion of additional metrics like pseudorange residuals as suggested in (Esswein & Psiaki, 2019), AGC or autocorrelation function should be explored. Further work will focus on leveraging the LRT-based detection framework as the foundation for such a combination.

ACKNOWLEDGEMENTS

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 United States government for providing us with an opportunity to test under live GPS spoofing. Special thanks go 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 conducting of the flight tests. The authors thank Professor J. David Powell for his invaluable advice and guidance during this research.

ORCID

Fabian Rothmaier  <https://orcid.org/0000-0002-9215-9881>

Sherman Lo  <https://orcid.org/0000-0002-4814-6506>

Todd Walter  <https://orcid.org/0000-0002-3257-3175>

REFERENCES

- Akos, D.M. (2012). Who's afraid of the spoofer? GPS/GNSS spoofing detection via automatic gain control (AGC). *NAVIGATION*, 59(4), 281–290. <https://doi.org/10.1002/navi.19>
- Appel, M., Iliopoulos, A., Fohlmeister, F., Pérez Marcos, E., Cuntz, M., Konovaltsev, A., ... Meurer, M. (2019). Experimental validation of GNSS repeater detection based on antenna arrays for maritime applications. *CEAS Space Journal*, 11(1), 7–19. <https://doi.org/10.1007/s12567-018-0232-6>
- Appel, M., Konovaltsev, A., & Meurer, M. (2015). Robust spoofing detection and mitigation based on direction of arrival estimation. *Proc. of the 28th International Technical Meeting of the Satellite Division of The Institute of Navigation*, Tampa, FL, 3335–3344.
- Bhatti, J. A., & Humphreys, T. E. (2017). Hostile control of ships via false GPS signals: Demonstration and detection. *NAVIGATION*, 64, 51–66. <https://doi.org/10.1002/navi.183>
- Blanch, J., Walter, T., Enge, P., Lee, Y., Pervan, B., Rippl, M., ... Kropp, V. (2014). Baseline advanced RAIM user algorithm and possible improvements. *IEEE Transactions on Aerospace and Electronic Systems*, 51(1), 713–732. <https://doi.org/10.1109/TAES.2014.130739>
- Borio, D., & Gioia, C. (2016). A sum-of-squares approach to GNSS spoofing detection. *IEEE Transactions on Aerospace and Electronic Systems*, 52(4), 1756–1768. <https://doi.org/10.1109/TAES.2016.150148>
- Chen, Y. H., Lo, S., Perkins, A., Rothmaier, F., Akos, D. M., & Enge, P. (2018). Demonstrating single element null steering antenna direction finding for interference detection. *Proc. of the 2018 International Technical Meeting of The Institute of Navigation*, Reston, VA, 240–259. <https://doi.org/10.33012/2018.15598>
- Chen, Y. H., Rothmaier, F., Akos, D., Lo, S., & Enge, P. (2017). Towards a practical single element null steering antenna, *Proc. of the 2017 International Technical Meeting of The Institute of Navigation*, Monterey, CA, 879–889. <https://doi.org/10.33012/2017.14954>
- Cook, S. A. (1971). The complexity of theorem-proving procedures. *Proc. of the Annual ACM Symposium on Theory of Computing*, 151–158. <https://doi.org/10.1145/800157.805047>
- Egea-Roca, D., Tripana-Caballero, A., López-Salcedo, J. A., Seco-Granados, G., De Wilde, W., Bougard, B., & Popugaev, A. (2018). GNSS measurement exclusion and weighting with a dual polarized antenna: The FANTASTIC project. *Proc. of the 2018 8th International Conference on Localization and GNSS*, Guimaraes, Portugal, 1–6. <https://doi.org/10.1109/ICL-GNSS.2018.8440897>

- Esswein, M. C., & Psiaki, M. L. (2019). GNSS anti-spoofing for a multi-element antenna array. *Proc. of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation*, Miami, FL, 3197–3214. <https://doi.org/10.33012/2019.17062>
- European Global Navigation Satellite Systems Agency. GNSS Market Report (2017). In *GNSS Market Report*. <https://doi.org/10.2878/0426>
- Fulton, C., Lee, C., Wright, D., Eastburg, G., Rivey, J., Whitney, S., & Atkins, T. (2020). Flight testing of advanced receiver autonomous integrity monitoring and dual polarized antenna. *2020 IEEE/ION Position, Location and Navigation Symposium, PLANS*, Portland, OR, 515–527. <https://doi.org/10.1109/PLANS46316.2020.9109903>
- Greenwood, D. T. (1987). *Principles of Dynamics*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall.
- Gross, J. N., Kilic, C., & Humphreys, T. E. (2019). Maximum-likelihood power-distortion monitoring for GNSS-Signal authentication. *IEEE Transactions on Aerospace and Electronic Systems*, 55(1), 469–475. <https://doi.org/10.1109/TAES.2018.2848318>
- Günther, C. (2014). A survey of spoofing and counter-measures. *NAVIGATION*, 61(3), 159–177. <https://doi.org/10.1002/navi.65>
- Humphreys, T. E., Ledvina, B. M., Psiaki, M. L., Hanlon, B. W. O., & Kintner, P. M. (2008). Assessing the spoofing threat: Development of a portable GPS civilian spoofer. *Proc. of the 21st International Technical Meeting of the Satellite Division of The Institute of Navigation*, Savannah, GA, 2314–2325.
- Hytti, H., & Visala, A. (2015). A DCM based attitude estimation algorithm for low-cost MEMS IMUs. *International Journal of Navigation and Observation*, 1–18. <https://doi.org/10.1155/2015/503814>
- Jafarnia-Jahromi, A., Broumandan, A., Nielsen, J., & Lachapelle, G. (2012). GPS vulnerability to spoofing threats and a review of antispoofing techniques. *International Journal of Navigation and Observation*, 1–16. <https://doi.org/10.1155/2012/127072>
- Konovaltsev, A., Caizzzone, S., Cuntz, M., & Meurer, M. (2014). Autonomous spoofing detection and mitigation with a miniaturized adaptive antenna array. *Proc. of the 27th International Technical Meeting of the Satellite Division of The Institute of Navigation*, Tampa, FL, 2853–2861.
- Konovaltsev, A., Cuntz, M., Haettich, C., & Meurer, M. (2013). Performance analysis of joint multi-antenna spoofing detection and attitude estimation. *Proc. of the 2013 International Technical Meeting of the Institute of Navigation*, San Diego, CA, 864–872.
- Lehmann, E. L., & Romano, J. P. (2005). *Testing Statistical Hypotheses* (3rd ed.). New York: Springer. <https://doi.org/10.1007/0-387-27605-X>
- Lo, S., Chen, Y. H., Jain, H., & Enge, P. (2018). Robust GNSS spoof detection using direction of arrival: Methods and practice. *Proc. of the 31st International Technical Meeting of the Satellite Division of the Institute of Navigation*, Miami, FL, 2891–2906. <https://doi.org/10.33012/2018.15900>
- Lo, S., Chen, Y. H., Rothmaier, F., Zhang, G., & Lee, C. (2020). Developing a dual polarization antenna (DPA) for high dynamic applications. *Proc. of the 2020 International Technical Meeting of the Institute of Navigation*, San Diego, CA, 1001–1020. <https://doi.org/10.33012/2020.17193>
- Magiera, J., & Katulski, R. (2015). Detection and mitigation of GPS spoofing based on antenna array processing. *Journal of Applied Research and Technology*, 13(1), 45–57. [https://doi.org/10.1016/S1665-6423\(15\)30004-3](https://doi.org/10.1016/S1665-6423(15)30004-3)
- Manfredini, E. G., Akos, D. M., Chen, Y. H., Lo, S., Walter, T., & Enge, P. (2018). Effective GPS spoofing detection utilizing metrics from commercial receivers. *Proc. of the 2018 International Technical Meeting of The Institute of Navigation*, Reston, VA, 672–689. <https://doi.org/10.33012/2018.15595>
- McMilin, E. (2016). *Single antenna null-steering for GPS & GNSS aerial applications* (Ph.D. dissertation). Stanford, CA: Stanford University.
- Meurer, M., Konovaltsev, A., Appel, M., & Cuntz, M. (2016). Direction-of-arrival assisted sequential spoofing detection and mitigation. *Proc. of the 2016 International Technical Meeting of The Institute of Navigation*, Monterey, CA, 181–192. <https://doi.org/10.33012/2016.13395>
- Meurer, M., Konovaltsev, A., Cuntz, M., & Hättich, C. (2012). Robust joint multi-antenna spoofing detection and attitude estimation using direction assisted multiple hypothesis RAIM. *Proceedings of the 25th Meeting of the Satellite Division of the Institute of Navigation*, Nashville, TN, 3007–3016.
- Pirsiavash, A., Broumandan, A., & Lachapelle, G. (2016). Two-dimensional signal quality monitoring for spoofing detection. Presented at Navitec, Noordwijk, Netherlands: ESA/ESTEC.
- Psiaki, M. L., & Humphreys, T. E. (2016). GNSS spoofing and detection. *Proc. of the IEEE*, 104(6), 1258–1270. <https://doi.org/10.1109/JPROC.2016.2526658>
- Psiaki, M. L., O'Hanlon, B. W., Powell, S. P., Bhatti, J. A., Humphreys, T. E., & Schofield, A. (2014). GNSS spoofing detection using two-antenna differential carrier phase. *Proc. of the 27th International Technical Meeting of the Satellite Division of The Institute of Navigation*, Tampa, FL, 2776–2800.
- Psiaki, M. L., Powell, S. P., & O'Hanlon, B. W. (2013). GNSS spoofing detection using high-frequency antenna motion and carrier-phase data. *Proc. of the 26th International Technical Meeting of the Satellite Division of the Institute of Navigation*, Nashville, TN, 2949–2991.
- Regional Aviation Safety Group for the Middle East Region (RASG-MID). (2019). Guidance material related to GNSS vulnerabilities. *RASG-MID Safety Advisory*, (14). ICAO. <https://www.icao.int/MID/Documents/2017/RASG-MID6/RSA%2014-GNSS%20Vulnerabilities.pdf>.
- Rothmaier, F., Chen, Y. H., Lo, S., & David Powell, J. (2019). Single GNSS antenna heading estimation. *Proc. of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation*, Miami, FL, 2159–2171. <https://doi.org/10.33012/2019.16915>
- Rothmaier, F., Chen, Y., & Lo, S. (2019). Improvements to steady state spoof detection with experimental validation using a dual polarization antenna. *Proc. of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation*, Miami, FL, 967–983. <https://doi.org/10.33012/2019.16989>
- Van Trees, H. L. (2001). *Detection, Estimation, and Modulation Theory, Part I*, New York: John Wiley & Sons.: <https://doi.org/10.1002/0471221082>
- Wesson, K. D., Gross, J. N., Humphreys, T. E., & Evans, B. L. (2018). GNSS signal authentication via power and distortion monitoring. *IEEE Transactions on Aerospace and Electronic Systems*, 54(2), 739–754. <https://doi.org/10.1109/TAES.2017.2765258>

How to cite this article: Rothmaier F, Chen Y-H, Lo S, Walter T. GNSS spoofing detection through spatial processing. *NAVIGATION*. 2021;68(2):243–258. <https://doi.org/10.1002/navi.420>