GNSS Spoofing, Jamming, and Multipath Interference Classification using a Maximum-Likelihood Multi-Tap Multipath Estimator

Jason N. Gross, *West Virginia University*
Todd E. Humphreys, *University of Texas at Austin*

BIOGRAPHY

Dr. Jason N. Gross is an assistant professor at West Virginia University (WVU) were he directs WVU's Navigation Laboratory. He received his Ph.D. in aerospace engineering from WVU in 2011 and the worked at the Jet Propulsion Laboratory (JPL) until January 2014. His research is focused on Guidance, Navigation & Control (GNC) technologies and he currently serves as an associate member of the AIAA's GNC technical committee. He is a member of the ION, the AIAA, and the IEEE.

Dr. Todd E. Humphreys is an associate professor in the department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin, and Director of the UT Radionavigation Laboratory. He received a B.S. and M.S. in Electrical and Computer Engineering from Utah State University and a Ph.D. in Aerospace Engineering from Cornell University. He specializes in the application of optimal detection and estimation techniques to problems in satellite navigation, autonomous systems, and signal processing.

ABSTRACT

This paper experimentally evaluates the application of existing multipath mitigation technology in conjunction with in-band power monitoring for the purpose of GNSS interference classification. In particular, GNSS interference detection and classification metrics derived from the output of a multiple-correlation tap, maximum-likelihood multipath estimator are jointly used for the alarming the presence of GNSS spoofing, jamming or multipath. This approach is evaluated against the Texas Spoofing Text Battery (TEXBAT) archive (Humphreys et al., 2012), a dozen sets of deep urban multipath recordings, several recordings of wideband jammers at several different power levels, and clean static data recordings. Two detection approaches are proposed and one is shown to be better at discriminating between spoofing and jamming attacks.

INTRODUCTION

The open-access standards of GNSS have made its use commonplace for within day-to-day civil applications worldwide. However, this openness has been shown to expose civil receivers to the vulnerability of being 'overtaken' by counterfeit spoofing signals that can then mislead the user to accepting a false position or time solution(Humphreys et al., 2008; Shepard et al., 2012; Kerns et al., 2014).

To defend against spoofing attacks, several algorithms have been developed and proposed. A comprehensive review of GNSS spoofing detection methods is provided in (Psiaki and Humphreys, 2016), in which the authentication methods used to defend against spoofing are broadly categorized into: (1) those that use cryptographic techniques, (2) those that use geometry and either multiple antennas to exploit angle-of-arrival checks, and (3) those that do not belong to (1) or (2).

The technique proposed and evaluated herein, falls into group (3), and is an extension of the so called 'pincer' approach that was introduced in (Wesson et al., 2013) and derived and evaluated in detail within (Wesson et al., 2016). The 'pincer' approach jointly monitors both correlation profile distortion with symmetric–difference measurements and received in-band power. The defense 'traps' a would-be attacker to a relatively low power-advantage so that the attack is not easily detectable by the power-monitor. This, in turn, yields increased level of signal distortion. The technique offers the important advantages that it is simple to implement, requires no additional hardware, and is incredibly effective at detecting the occurrence of a malicious attack (i.e., spoofing or jamming) from data that is only interfered by natural multipath, however, once a receiver as been captured by a spoofer, a significant percentage of spoofed data is classified as a jamming attack. Therefore, in an effort

to better discriminate between a spoofing attack and a jamming attack, this paper's contribution is extends upon the 'pincer' concept by investigating the use of a multi-tap maximum likelihood multipath estimator as the distortion monitor.

The rest of this paper is organized as follows. Section 1 discuses the proposed approach and includes of a review of the particular multipath estimator employed, the model adopted for the GNSS signal with a single interference source, the simulation of the observation space of the proposed detector, and the application of a multi–hypothesis Bayesian classifier for interference classification. Section 2 introduces the data sets used for experimental evaluation and presents the proposed method's classification performance. Finally concluding remarks, including next steps for this work are offered in Section 3.

1 TECHNICAL APPROACH

To monitor correlation profile distortion, a multi-correlation tap, maximum likelihood multipath estimator as detailed in (Blanco-Delgado and Nunes, 2012) is used decompose the raw GNSS In-Phase and Quadrature samples into two estimated signal models. From these estimated signal models, the interference detection and classification metric are defined using the amplitude of the two estimated correlation profiles based upon the following simple intuition:

1. Data that are subjected to only thermal noise will yield a significant amplitude on only one of the two estimated signal models;
2. Data that subjected to a GNSS spoofing attack will yield significant amplitudes on both estimated signal models (i.e., both the authentic and counterfeit GNSS signals are valid and yield a significant correlation amplitude);
3. Data that are subjected to jamming interference will not yield a amplitude for either of the estimated signal models;
4. Data that are subjected to multipath interference will look similar to spoofed data with the exception that the amplitude of correlation profile of the decomposed non line-of-sight signal reflection s typically smaller then a spoofed signal.

Next, the multipath estimator used in this study is reviewed.

1.1 Multi-Tap Maximum-Likelihood Multipath Estimation

This study assumes access to a receiver that employs multiple correlation taps and employs the maximum likelihood multipath estimation technique that leverages multiple tap receivers. The estimator and its detailed recipe for implementation is provided in (Blanco-Delgado and Nunes, 2012), however, its basic premise and structure are reviewed here for completeness. In particular, the multiple taps of post-correlation IQ data, U_k are decomposed into two models of the following form

$$U_k = \sum_{i=0}^{1} a_i e^{j\phi_i} \tilde{R}(\delta_k - \tau_i) + N_k \tag{1}$$

where, a_i is i^{th} signal the amplitude, i^{th} signal ϕ_i is the code phase, $\tilde{R}$ is the GPS cross-correlation function, δ_k is the k^{th} tap offset, τ_i is the i^{th} signal delay, and N_k is thermal noise, which has the following linear form:

$$U = H(\tau)\theta(a,\phi) \tag{2}$$

where $H(\tau)$ only is only dependent on the unknown delays and the GPS correlation function $\tilde{R}$, and for the case of $+/-N$ taps and two estimated signals is given as (Blanco-Delgado and Nunes, 2012)

$$H(\tau) = \begin{bmatrix} \tilde{R}(\delta_{-N} - \tau_0) & \dots & \tilde{R}(\delta_{-N} - \tau_1) \\ \vdots & \vdots & \vdots \\ \tilde{R}(\delta_{-N}\tau_0) & \dots & \tilde{R}(\delta_N - \tau_1) \end{bmatrix}. \tag{3}$$

Likewise, the estimated amplitude and phase is given as $\theta(a,\phi) = \begin{bmatrix} a_0 e^{j\phi_0} & a_1 e^{j\phi_1} \end{bmatrix}$

The key of the estimator is to exploit the availability of multiple taps at known delay spacings in order to separate estimation of the $\tau_i's$ from $\theta(a,\phi)$ by considering all possible delay combinations available amongst the available signal taps. For example, for the proposed application, where the IQ data is decomposed into two signals, the result is a 2-D grid of possible delay combinations that are each assumed, and a corresponding $\theta(a,\phi)$ is estimated. When fixing τ to a grid point, the estimation problem then become s

$$\hat{\theta} = (H^T Q^{-1} H) H^T Q^{-1} U \tag{4}$$

where Q is the Toeplitz matrix that accounts for the complex Gaussian correlated noise of the multi-tap IQ samples and is defined as $Q_{i,j} = \tilde{R}(|i-j|\Delta)$ and Δ is the tap spacing (Blanco-Delgado and Nunes, 2012). After testing all delay combinations on the grid, the delay combinations for the solution with the smallest squared error residual $|U - H\theta|_{Q^{-1}}$ is accepted.

Furthermore, it is important to note that for classic multipath estimation, the evaluated delay combination grid points can be further pruned to account for the fact that all reflected signals must have a that is delay larger than the line-of-sight signal (i.e., $\tau_1 > \tau_0$), however, for the application proposed herein, this simplification cannot be made, as there is no such restriction on a spoofer's counterfeit signal.

Additionally, It is worth noting, as pointed out by Blanco-Delgado and Nunes (2012), that for each possible delay combination, the matrix, $(H^T Q^{-1} H) H^T Q^{-1}$, is completely dependent on the receiver tap configuration and the GNSS cross-correlation function, and can therefore be precomputed offline and stored, thus drastically increasing the real-time performance of this estimator. In addition, for more resolution, an interpolation scheme can be employed to estimated delays the fall between tap offsets, however, this feature was excluded in the present study.

1.2 Proposed Classification Observation and Observation Modeling

Based on the output of the above referenced multipath estimator, based on the high-level intuition enumerated in the introduction, two observations were considered for monitoring correlation distortion, namely, the estimated amplitude of the second reflection, $\hat{a_1}$, and the summation of the two estimated amplitudes $\hat{a_0} + \hat{a_1}$. That is, a significant $\hat{a_1}$ should only be present under spoofing and multipath, and furthermore, the use of $\hat{a_0} + \hat{a_1}$ is expected to reveal that neither $\hat{a_0}$ or $\hat{a_1}$ is significant under jamming.

In order to characterize the expected distribution of these metrics, a post-correlation model of GPS with a single interference source was adopted. In particular, the model from (Van Nee, 1993) that is detailed and adopted in (Wesson et al., 2016) was employed.

$$\xi_k(\tau) = \beta_k[\xi_{\mathrm{A}k}(\tau) + \xi_{\mathrm{I}k}(\tau) + \xi_{\mathrm{N}k}(\tau)] \tag{5}$$

where β_k is the average value of the Automatic Gain Control (AGC) scaling $\beta(t)$ over the kth accumulation interval and $\xi_{\mathrm{A}k}(\tau)$, $\xi_{\mathrm{I}k}(\tau)$, $\xi_{\mathrm{N}k}(\tau)$ are the complex correlation function components corresponding to the authentic signal, the interference signal, and thermal noise, respectively.

The correlation components $\xi_{\mathrm{A}k}(\tau)$ and $\xi_{\mathrm{I}k}(\tau)$ can be modeled in terms of GPS auto-correlation function $\tilde{R}(\tau)$ as

$$\xi_{\mathrm{A}k}(\tau) = \sqrt{P_{\mathrm{A}k}} R(-\Delta\tau_{\mathrm{A}k} + \tau) \exp(j\Delta\theta_{\mathrm{A}k}) \tag{6}$$

$$\xi_{\mathrm{I}k}(\tau) = \sqrt{\eta_k P_{\mathrm{A}k}} R(-\Delta\tau_{\mathrm{I}k} + \tau) \exp(j\Delta\theta_{\mathrm{I}k}) \tag{7}$$

where $P_{\mathrm{A}k}$ and η_k are the average values of the authentic signal power P_{A} and η over the accumulation interval, and $\Delta\tau_{\mathrm{A}k}$ is the average value of the difference $\tau_{\mathrm{A}} - \hat{\tau}$ over the accumulation interval, with similar definitions for $\Delta\tau_{\mathrm{I}k}$, $\Delta\theta_{\mathrm{A}k}$, and $\Delta\theta_{\mathrm{I}k}$. The assumed distribution of these parameters, under each Hypothesis, $H0$-null, $H1$-multipath, $H2$-spoofing, $H3$-jamming are described next.

As introduced in (Wesson et al., 2016), three key parameters, denoted as Θ, within the above model are relevant to distinguishing spoofing, jamming and multipath. First, η, which is the power advantage of the interference signal. Next, $\Delta\tau \equiv \tau_{\mathrm{I}} - \tau_{\mathrm{A}}$, and finally $\Delta\theta \equiv \theta_{\mathrm{I}} - \theta_{\mathrm{A}}$, which are the relative delay and phase between the interfered and the authentic signal. These three parameters make up Θ, which can be assumed to model each interference type

1.2.1 Parameters under H_0: Thermal Noise Only

Under the hypothesis H_0 that only thermal noise is present, $\Theta = [0, 0, 0]$.

1.3 Parameters under H_1: Multipath Interference

Under the hypothesis that the interference is due to multipath, H_1, each element in θ can be modeled by a different probability distribution. In this study, we assume that $\eta \sim$ Rayleigh, $\tau \sim$ Exponential, and $\phi \sim$ Uniform.

To select the distribution parameters used for each element of θ, bounds were selected for each element of θ based upon reasonable assumptions for mulitpath. In particular,

$$\mathcal{X}_1 = \{\Theta \in \mathcal{X} | 0 < \eta \leq \eta_1, 0 < \Delta\tau < 2\tau_c, 0 < \theta < 2\pi]\} \tag{8}$$

where $\eta_1 = 0.7$ was selected under the assumption that a multipath reflection has a very low probability actually capturing the tracking loops. This led to $\eta \sim \text{Rayleigh}(\sigma_{\mathrm{r}} = 0.29)$. For τ, it is know that a multipath reflection must be delayed with respect to the authentic signal, such that $\tau > 0$, and upper bound of $2\tau_c$ was selected. Finally, $\theta_{AI} \sim U(0, 2\pi)$ was assumed.

1.3.1 Parameters under H_2: Spoofing Interference

The hypothesis H_2 is that the interference is due to a spoofer. In this case, the spoofer could conceivably mimic or select any varied selection of θ. Thus, the model must make broad and general assumptions, and as such, parameter bounds under spoofing were chose as

$$\mathcal{X}_2 = \{\Theta \in \mathcal{X} | \eta_1 < \eta \leq \eta_2, -2\tau_c < \Delta\tau < 2\tau_c, 0 < \theta < 2\pi]\} \tag{9}$$

First, for modeling the distribution of η, assume that the spoofer's motive is to capture the receiver without being blatantly obvious with respect an in-band power monitor. That this, while a spoofer can conceivably select any power advantage, a reasonable assumption is to design around the worst-case of a near power-matched spoofing. Therefore, the distribution of η was modeled Rayleigh, just as the multipath, but with lower of bound chosen as the upper bound of the multipath case (i.e., η_1). Furthermore, the spread of the distribution was chosen to be relatively narrow, so as to defend against power-matched spoofing with more emphasis than large power advantage attacks. As such, $\eta \sim \eta_1 + \text{Rayleigh}(\sigma_r = 6)$.

Second, assume that at the onset of the attack, in order to minimize the distortion caused the attach, the attacker attempts to minimize τ. However, once the tracking loops have been commandeered, unlike multipath, τ can be positive or negative. Therefore, $\tau \sim N(0, 1\tau_c)$ was chosen to reflect these assumptions. $\theta_{AI} \sim U(0, 2\pi)$ was also assumed for H_2.

1.3.2 Parameters under H_3: Jamming Interference

Under the hypothesis H_3, the interference is due to a jammer. Parameters were assumed to be bounded as

$$\mathcal{X}_2 = \{\Theta \in \mathcal{X} | \eta_1 < \eta \leq \eta_3, |\Delta\tau| > 2\tau_c, 0 < \theta < 2\pi]\}. \tag{10}$$

Just as spoofing, a lower bound of η_1 was assumed, however, the spread of η was assumed to be wider for jamming. The power advantage under H_3 was therefore assumed to be given as $\eta \sim \eta_1 +$

Rayleigh($\sigma_r = 12$).

Further, a uniform distribution with $|\Delta\tau| > 2\tau_c$ was chosen to reflect the fact that a jamming signal will not yield an interference source that has GNSS correlation profile. Again $\theta_{AI} \sim U(0, 2\pi)$ was also assumed of H_3.

1.3.3 Simulated Observation Space

Using the assumed parameter distributions and the post-correlation model described above, a Monte-Carlo was conducted in which 500,000 signals were simulated. For each signal, 41 IQ taps were simulated and fed into the multipath estimator outlined in Section 1.1 and the amplitudes were then normalized to account for AGC scaling. In addition, the total in-band power measurements were simulated as discussed in (Wesson et al., 2016), such that the total detector observations that were simulated $z_{0,k} = \left[\hat{a_1} P_k\right]$ and $z_{1,k} = \left[\hat{a_0} + \hat{a_1} P_k\right]$. The simulated observations space for z_0 and z_1 are shown in Figures 1 and 2.

As shown in Figures 1 $\hat{a_1}$, as expected seems to be a good proxy for distortion, as the z_0 observation space looks nearly identical to the symmetric-difference distortion metrics using in (Wesson et al., 2013, 2016). In addition, within Figure 2, at high-powers, the separation between jamming and spoofing look promising for their using in better distinguishing between these attacks.

1.4 Bayesian Classifier

Based on the output of the Monte-Carlo simulation, of there observation space, 2-D kernel density estimation was employed to numerically realize approximations for the probability transition mechanism for each $p(z_0|\Theta_{0,1,2,3})$ and $p(z_1|\Theta_{0,1,2,3})$.

With modeled distributions for each H0—H3, a Bayesian framework of probability and cost is applicable to classify the anomalies while minimizing risk. Let the notation C_{ij} indicate that i is declared when j is true. Herein, the simplest approach and test the classifier is adopted for penalizing completely for misclassification error. To realize this, a cost matrix ,C, was selected to be

$$C = \begin{bmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{bmatrix} \tag{11}$$

Furthermore, assumed prior probabilities that the null hypothesis (i.e., thermal noise) is most likely, $p_0 = 0.4$, followed by multipath, $p_1 = 0.3$, spoofing, $p_2 = 0.2$, and jamming, $p_3 = 0.1$ were used, such that the optimal classification, decision at each

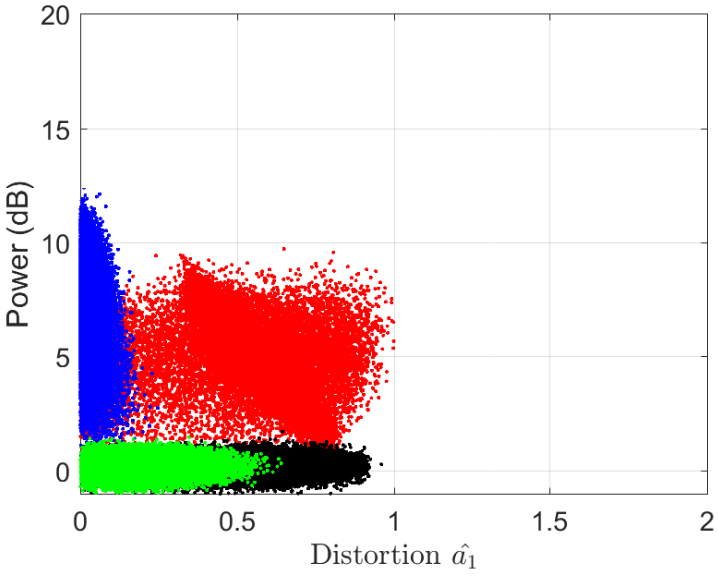

Figure 1. Scatter plot showing simulated z_0 for clean (green), multipath (black), spoofing (red), and jamming (blue) using the probability distribution models developed in the previous section.

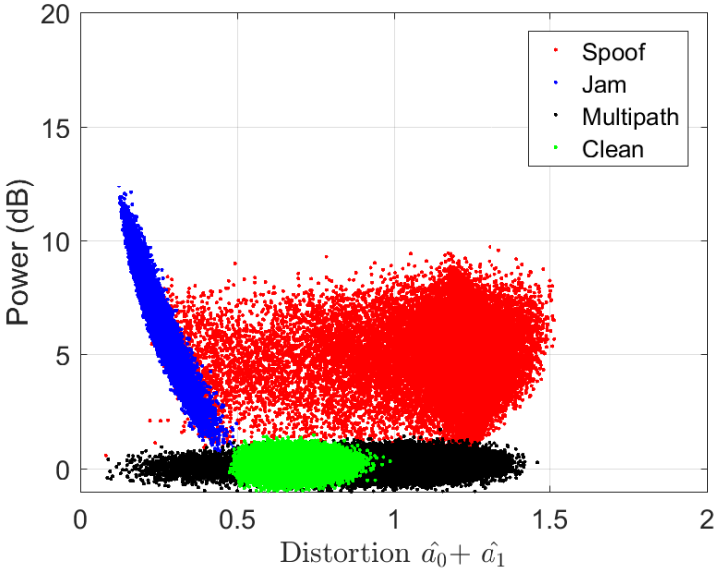

Figure 2. Scatter plot showing simulated z_1 for clean (green), multipath (black), spoofing (red), and jamming (blue) using the probability distribution models developed in the previous section.

grid for each possible z_0 or z_1 was determined as follows

$$\Gamma_k = C \begin{bmatrix} p_0 p(z_k|\Theta_0) \\ p_1 p(z_k|\Theta_1) \\ p_2 p(z_k|\Theta_2) \\ p_3 p(z_k|\Theta_3) \end{bmatrix} \tag{12}$$

where z_k is either $z_{0,k}$ or $z_{1,k}$ for all grid points k, Γ_k is the risk of each classification at grid point k, and the decision associated in argmin(Γ_k) is optimal.

Following this procedure, the decision revisions for z_0 and z_1 were determined to be as shown in Figures 3 and 4

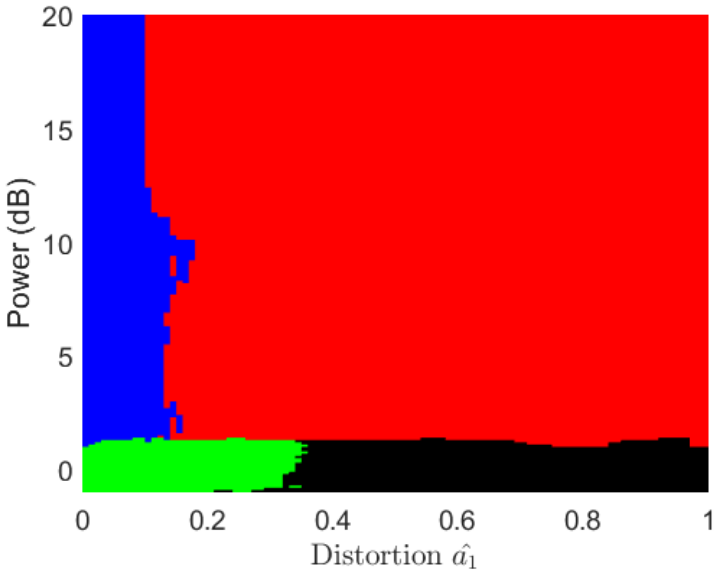

Figure 3. Optimal decision regions to minimized misclassification with z_0 for clean (green), multipath (black), spoofing (red), and jamming (blue) using the probability distribution models developed in the previous section.

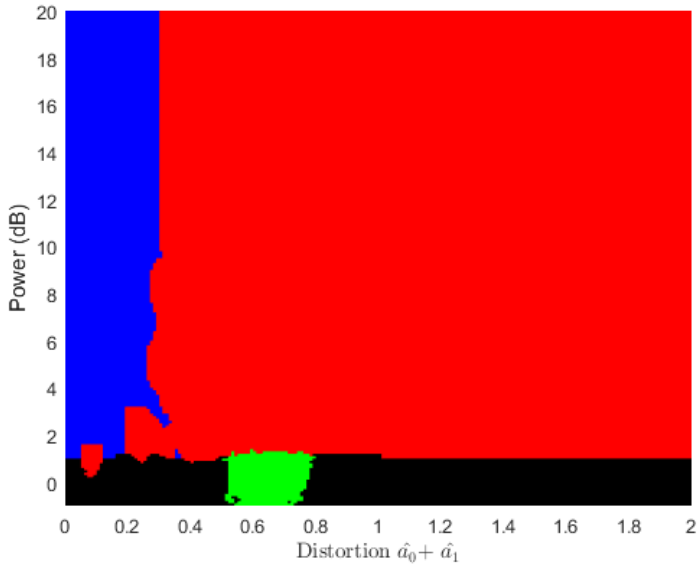

Figure 4. Optimal decision regions to minimized misclassification with z_1 for clean (green), multipath (black), spoofing (red), and jamming (blue) using the probability distribution models developed in the previous section.

2 RESULTS

2.1 Experimental Data

To evaluate the proposed technique the same experimental data recordings that were used in (Wesson et al., 2016) were used for an independent experimental evaluation of the proposed classification methods. In total, 26 data recordings were used for evaluation. As detailed in the following subsections, these included 6 of the TXBAT spoofing , 15 recordings of live-sky multipath-dense scenarios, 4 jamming scenarios of different power levels, and 1 live-sky scenarios subjected to only thermal noise.

2.1.1 TEXBAT 1.1

The Texas Spoofing Test Battery (TEXBAT) 1.1 is a set of eight high-fidelity digital recordings of spoofing attacks against the civil GPS L1 C/A signals (Humphreys et al., 2012). Both stationary- and dynamic-receiver-platform scenarios are provided along with their corresponding un-spoofed recording. Each 16-bit quantized recording was centered at the GPS L1 C/A center frequency of 1575.42 MHz with a bandwidth of 20 MHz and at a complex sampling rate of 25 MSps.

2.1.2 RNL Multipath and Interference Recordings

Additional data was collected where the receiver was statically positioned in multipath environments or recordings while driving in multipath-dense environments (e.g., urban canyon). Each 16-bit quantized recording was centered at the GPS L1 C/A center frequency with a bandwidth of 10 MHz and at a complex sampling rate of 12.5 MSps.

2.1.3 Jamming

Jamming noise was recorded from a "cigarette lighter jammer" with a sweep range of 1550.02–1606.72 MHz and sweep period of 26 μ. The recorded jamming noise was combined with clean, static-receiver data from a rooftop antenna and re-recorded. Each 16-bit quantized recording was centered at the GPS L1 C/A center frequency with a bandwidth of 10 MHz and at a complex sampling rate of 12.5 MSps.

2.1.4 Thermal Noise Only

A recording under static rooftop and open-field environments included the un-spoofed recording within the TEXBAT data sets.

2.2 Evaluation

To provide for a simple evaluation, all data were pre-marked as either H0-clean, H1-multipath, H2-spoofing, H3-jamming. For spoofing, jamming and clean, this process was straight forward, as the attack scenarios were known, however, for mulitpath, much of the urban and dynamic recordings were largely clean and only a small portion were actually experiencing multipath interference. As such, for the case of multipath, the data were pre-sorted using the empirically estimated probability density for clean data, in order to classify much of the interference recordings as clean. The recorded observation space for both z_0 and z_1 are shown in Figures 5 and 6.

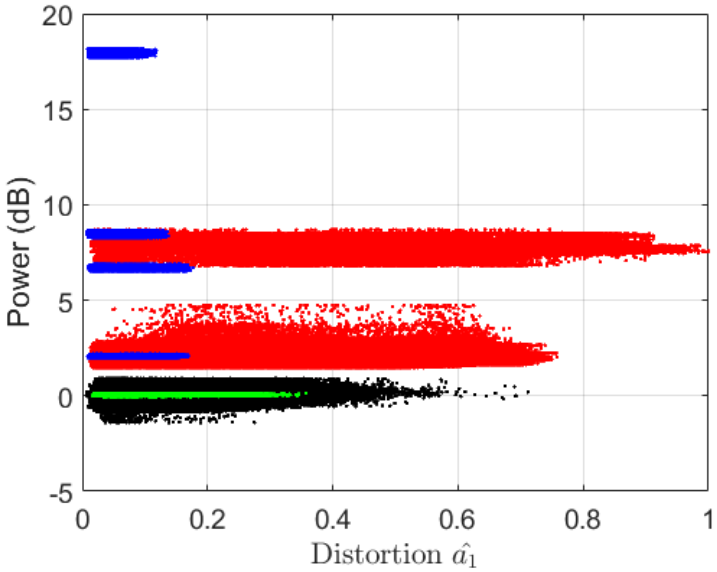

Figure 5. Optimal decision regions to minimized misclassification with z_0 for clean (green), multipath (black), spoofing (red), and jamming (blue) using the probability distribution models developed in the previous section.

Finally, to evaluate the classifiers, the recorded observation Figures 5 and 6 were classified using the decision regions derived from the Monte-Carlo simulation shown in Figures 3 and 4. Tables 1 and 2 show the classification rates.

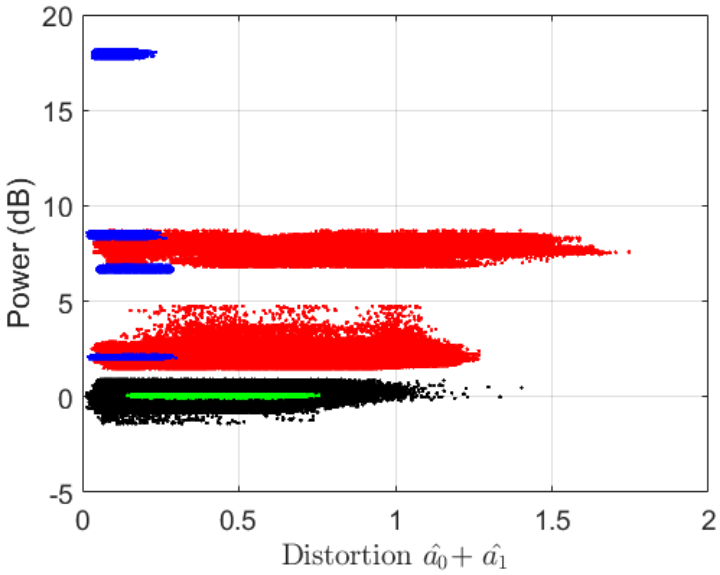

Figure 6. Optimal decision regions to minimized misclassification with z_1 for clean (green), multipath (black), spoofing (red), and jamming (blue) using the probability distribution models developed in the previous section.

Table 1. Classification rates when using $z_0 = \begin{bmatrix}\hat{a}_1 & P\end{bmatrix}$ and evaluating against the experimental data recording.

	Decision			
	H_0	H_1	H_2	H_3
$H0$	0.999	.001	0 .0	0.0
$H1$	0.272	0.724	0.0	0.0
$H2$	0.0	0.0	0.727	0.273
$H3$	0.0	0.0	0.0131	0.987

As shown in Tables 1 and 2, using $\hat{a}_1$, much like using symmetric differences, has very little false positives as nearly all the clean data is classified as such, however, nearly 30% of the spoofing data is labeled as jamming. On the contrary, with the proposed $z_1 = \begin{bmatrix}\hat{a}_0 + \hat{a}_1 & P\end{bmatrix}$, a severe drop in performance occurs with respect to multipath false positives, however, there is a nearly 20% improvement with respect to the discrimination between spoofing and jamming.

3 CONCLUSIONS

A GNSS interference classification method that builds upon the 'pincer' defense, but instead uses a maximum likelihood multipath estimator to monitor correlation profile distortion has been presented. When using $\hat{a}_1$ estimates alone for distortion, very similar performance is obtained as with the original 'pincer' that uses symmetric difference measurements. When combining $\hat{a}_0 + \hat{a}_1$ improved discrimination between spoofing and jamming has been shown to the level of 20%, however, this comes at the expense of many more mulitpath false positives when data is clean. Therefore, depending on the environment (i.e., if one is known to be in contested environment vs. a typical civil application) one approach may be favorable over the other. Future work will include: the development of a detection method that jointly uses both proposed metrics in conjunction, a direct comparison

Table 2. Classification rates when using $z_1 = \begin{bmatrix}\hat{a}_0 + \hat{a}_1 & P\end{bmatrix}$ and evaluating against the experimental data recording.

	Decision			
	H_0	H_1	H_2	H_3
$H0$	0.252	.748	0 .0	0.0
$H1$	0.0	0.978	0.022	0.0
$H2$	0.0	0.0	0.888	0.111
$H3$	0.0	0.0	0.069	0.931

with the pincer symmetric difference distortion monitor, and an evaluation the sensitivity of the proposed approaches to much fewer than 41 correlation taps (e.g., 7 taps are used in some commercial receivers).

ACKNOWLEDGEMENTS

J. Gross was supported for this work in part by the West Virginia University Big XII Faculty Fellowship Program

REFERENCES

N. Blanco-Delgado and F. D. Nunes. Multipath estimation in multicorrelator gnss receivers using the maximum likelihood principle. *IEEE Transactions on Aerospace and Electronic Systems*, 48(4):3222–3233, 2012.

T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. OHanlon, and P. M. Kintner Jr. Assessing the spoofing threat: Development of a portable gps civilian spoofer. In *Proceedings of the ION GNSS international technical meeting of the satellite division*, volume 55, page 56, 2008.

T. E. Humphreys, J. A. Bhatti, D. P. Shepard, and K. D. Wesson. The Texas Spoofing Test Battery: Toward a standard for evaluating GNSS signal authentication techniques. In *Proceedings of the ION GNSS Meeting*, 2012. `http://radionavlab.ae.utexas.edu/texbat`.

A. J. Kerns, D. P. Shepard, J. A. Bhatti, and T. E. Humphreys. Unmanned aircraft capture and control via gps spoofing. *Journal of Field Robotics*, 31(4):617–636, 2014.

M. L. Psiaki and T. E. Humphreys. Gnss spoofing and detection. *Proceedings of the IEEE*, 104(6):1258–1270, 2016.

D. P. Shepard, T. E. Humphreys, and A. A. Fansler. Evaluation of the vulnerability of phasor measurement units to gps spoofing attacks. *International Journal of Critical Infrastructure Protection*, 5(3):146–153, 2012.

R. D. Van Nee. Spread-spectrum code and carrier synchronization errors caused by multipath and interference. *IEEE Transactions on Aerospace and Electronic Systems*, 29(4):1359–1365, 1993.

K. Wesson, J. N. Gross, B. L. Evans, and T. E. Humphreys. Gnss signal authentication via joint detection of correlation function distortion and anomalous received power. *Submitted to IEEE Transactions on Aerospace and Electronic Systems*, 2016.

K. D. Wesson, B. L. Evans, and T. E. Humphreys. A combined symmetric difference and power monitoring gnss anti-spoofing technique. In *Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE*, pages 217–220. IEEE, 2013.