A Testbed for Developing and Evaluating GNSS Signal Authentication Techniques

Todd Humphreys, Jahshan Bhatti, Daniel Shepard, and Kyle Wesson

***Abstract*—An experimental testbed has been created for developing and evaluating Global Navigation Satellite System (GNSS) signal authentication techniques. The testbed advances the state of the art in GNSS signal authentication by subjecting candidate techniques to the strongest publicly-acknowledged GNSS spoofing attacks. The testbed consists of a real-time phase-coherent GNSS signal simulator that acts as spoofer, a real-time software-defined GNSS receiver that plays the role of defender, and post-processing versions of both the spoofer and defender. Two recently-proposed authentication techniques are analytically and experimentally evaluated: (1) a defense based on anomalous received power in a GNSS band, and (2) a cryptographic defense against estimation-and-replay-type spoofing attacks. The evaluation reveals weaknesses in both techniques; nonetheless, both significantly complicate a successful GNSS spoofing attack.**

Keywords: Cryptographic signal authentication, GNSS security, GNSS spoofing detection.

I. Introduction

Authentication of civil Global Navigation Satellite System (GNSS) signals is increasingly a concern. Spoofing attacks, in which counterfeit GNSS signals are generated for the purpose of manipulating a target receiver's reported position or time, have been demonstrated with low-cost commercial equipment against a wide variety of civil Global Positioning System (GPS) receivers [1]–[3]. Such attacks threaten the security of financial transactions, communications, power distribution, and transportation, which all depend on GNSS signals for accurate positioning and timing [4]–[8].

Whereas the military GPS waveform is by design unpredictable and therefore resistant to spoofing [9], civil GPS waveforms—and those of other civil GNSS—are unencrypted, unauthenticated, and openly specified in publicly-available documents [10], [11]. Also, although not entirely constrained by the signal specifications, the navigation data messages modulating these civil waveforms are highly predictable. The combination of known signal structure and data bit predictability makes civil GNSS signals an easy target for spoofing attacks.

A number of promising methods are currently being developed to defend against civil GNSS spoofing attacks. These can be categorized as (1) receiver-autonomous signal-processing-based techniques, which require no antenna motion or specialized hardware apart from the GNSS receiver itself [12]–[18]; (2) receiver-autonomous antenna-based techniques, which require antenna motion or specialized antenna hardware [19]–[25]; (3) receiver-autonomous techniques based on fusing GNSS observables with measurements from non-GNSS sensors such as inertial sensors [26]; (4) cryptographic techniques that require signal specification modifications to overlay unpredictable but verifiable modulation on existing or future civil GNSS signals [27]–[29]; and (5) techniques that exploit the existing encrypted military signals to offer civil GPS signal authentication for networked GPS receivers [30]–[33]. The best protection against GNSS spoofing will likely involve a combination of these.

Existing and proposed GNSS signal authentication schemes are all premised on hypothesis tests involving statistical models for the authentic and counterfeit GNSS signals. In general, the statistics of the null hypothesis (only authentic signals present) are well known and readily verified by laboratory experiment, but the statistics of the alternative hypothesis (spoofing attack underway) are poorly characterized, for two reasons. First, the exact parameters of a spoofing attack (e.g., spoofing signal power, number of spoofing signal transmitters, initial spoofing signal code and carrier phase alignment with authentic signals, etc.) are typically unknown to a defender; at best a defender can assume only an approximate probability distribution for such parameters. Second, in constructing a model to describe the alternative hypothesis, one often makes simplifying assumptions to facilitate analytical treatment of the detection problem. Thus, even if the spoofing parameters were perfectly known, the modeled distribution and the true distribution may differ in important ways.

The uncertainty involved in characterizing the alternative hypothesis points to the need for model validation via experiment. Unfortunately, GNSS signal generation hardware capable of the most sophisticated spoofing attacks is neither commercially available nor straightforward to construct. Thus, for example, experimental validation of the authentication technique proposed in [30] was limited to the null hypothesis, and validation of the technique proposed in [15] was limited to an unsophisticated repeater-spoofer attack scenario, which, as will be shown herein, led to an overly optimistic performance assessment. A testbed capable of simulating sophisticated and realistic spoofing attacks is needed so that the efficacy of proposed GNSS signal authentication techniques can be experimentally evaluated.

This paper makes two primary contributions. First, it describes an experimental testbed that has been created for developing and evaluating GNSS signal authentication techniques. The testbed consists of a software-defined real-time

Authors' addresses: Todd Humphreys, Jahshan Bhatti, and Daniel Shepard, Department of Aerospace Engineering, The University of Texas at Austin, Austin TX, 78712, Email: (todd.humphreys@mail.utexas.edu), (jahshan@utexas.edu), (dshepard.ut@gmail.com). Kyle Wesson Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin TX, 78712, Email: (kyle.wesson@utexas.edu).

phase-coherent GNSS signal simulator capable of carrying out sophisticated spoofing attacks, a real-time software-defined GNSS receiver that plays the role of defender, and post-processing versions of both the spoofer and defender. Previous work has exercised the testbed or its spoofer component [2], [3], [29], [32], [34]–[36], but this paper is the first to describe the testbed as such and to offer a comprehensive view of its capabilities. The paper's second primary contribution is an analytical and experimental evaluation of two recently-proposed civil GNSS signal authentication techniques, the received power spoofing detector proposed in [15] and the security code estimation and replay (SCER) attack defense proposed in [29]. In the course of evaluating the SCER attack defense, the paper details how the defense can be implemented in practice within a GNSS receiver. This will be useful for receiver manufacturers in the event that proposed techniques for modulating cryptographic signatures on broadcast civil GNSS signals get implemented [28], [37].

The following section describes the testbed. Thereafter, the two signal authentication schemes are introduced and evaluated.

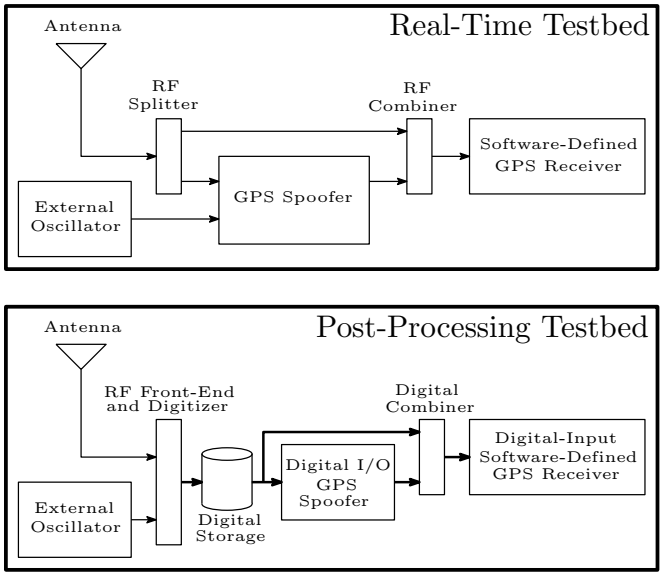

Fig. 1. Schematics depicting the real-time and post-processing versions of the signal authentication testbed. Thin lines in both schematics represent coaxial cables conveying analog signals, whereas thick lines in the lower schematic represent digital data streams.

II. Testbed Description

The real-time version of the signal authentication testbed consists of an advanced version of the GPS L1 C/A spoofer originally presented in [1] and a real-time software-defined GNSS receiver that plays the role of defender. A post-processing version of the testbed has also been developed to allow more flexibility in iterated testing of various spoofer and defender strategies. Schematics of both versions are shown in Fig. 1.

In the real-time testbed, the spoofer ingests authentic GPS L1 radio-frequency (RF) signals and outputs a counterfeit GPS L1 C/A RF signal ensemble. The counterfeit ensemble is combined with the original authentic signal ensemble in an RF combiner, and the composite authentic-counterfeit ensemble is directed to a software-defined GPS receiver. The spoofer can operate using its internal temperature-compensated crystal oscillator, but is most often driven by a higher-quality external oscillator to ensure minimal apparent variation in the time solution implied by the counterfeit signal ensemble.

In the post-processing testbed all signal processing downstream of the RF front-end operates on digital samples instead of analog RF signals. The end-to-end processing sequence is as follows: (1) the incoming authentic GPS L1 signals are digitized and stored, (2) the spoofer ingests the stored digital signals and outputs signals in a digital form, (3) the spoofer's output signals are combined with the digitized authentic signal stream via sample-wise digital multiplexing, (4) the receiver operates directly on the multiplexed digital data.

A. Spoofer

The University of Texas GPS spoofing device, shown in Fig. 2, is an advanced version of the original spoofer introduced in [1]. To the authors' knowledge, it is the most sophisticated publicly-acknowledged spoofing device. The latest version is capable of simultaneously tracking and spoofing up to 14 GPS L1 C/A signals while continuously attempting to acquire emerging GPS satellite signals. Other key features of the spoofer relevant to the testbed are phase alignment, navigation data bit prediction, variable output attenuation, noise padding, arbitrary generation of parity-correct navigation data streams, and SCER attack capability.

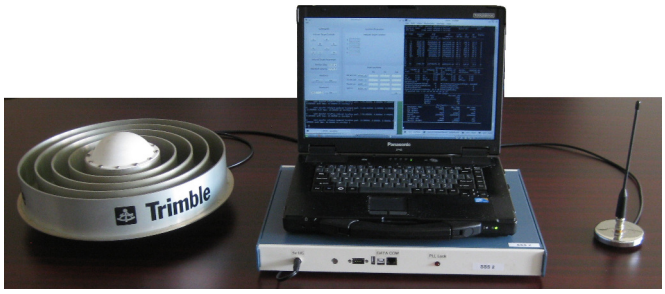

Fig. 2. The University of Texas real-time GPS spoofing device. The device as shown here is configured for over-the-air transmission but is only used as such in authorized tests [2], [3], [38]. More commonly, the spoofer's output signals are conveyed to the target receiver by coaxial cable or digital data. The computer shown atop the spoofing device runs a client application that allows a user to monitor and control a spoofing attack over a wireless or wireline network. The client-spoofer network connection is insensitive to latencies of hundreds of milliseconds, which permits a spoofing attack to be controlled remotely over the Internet.

1) Phase Alignment: The spoofer receives authentic civil GPS L1 C/A and GPS L2C signals and generates counterfeit GPS L1 C/A signals that are code-phase aligned with their authentic counterpart signals to within a few nanoseconds. In the real-time testbed, code-phase alignment is achieved by signal feedback. During a post-turn-on calibration phase, the spoofer acquires and achieves phase and data lock on all available authentic GPS L1 C/A signals. It then generates a simulated RF GPS signal whose spreading code is different from those modulating each of the authentic signals. It feeds

this unique signal back from its RF output to its RF input via an internal RF switch. At this point, the spoofer is able to acquire and track its own feedback signal in addition to the available authentic signals. By measuring the average offset between the feedback signal's received and transmitted code phase over an interval of time, the spoofer is able to precisely determine its own digital and analog latency. The latency, which amounts to approximately 5 ms, varies from turn-on to turn-on but remains constant to within the measurement precision thereafter, as can be verified by repeated calibration.

In the post-processing testbed, there is no need to compensate for processing latency, but the output of the digital input/output (I/O) spoofer must nonetheless be nanosecond-aligned with the digitized authentic signal stream. This is effected by sample-level adjustment in the digital combiner and sub-sample-level adjustment in the digital I/O spoofer.

As the spoofer attempts to induce a position or timing deviation in the target receiver by shifting the code phase of its counterfeit signals, it can adopt either of two strategies with respect to carrier phase generation. In the default mode, the rate of change of its signals' carrier phase is proportional to the rate of change of the corresponding code phase. Let $\dot{\tau}$ and $\dot{\theta}$ represent the rate of change of code phase and carrier phase, in seconds per second and radians per second, respectively. Then in the spoofer's default mode these are related by

$$\dot{\theta} = 2\pi f_c \dot{\tau} \tag{1}$$

where f_c is the GPS L1 frequency in Hz.

In an alternative mode, the so-called frequency lock mode, the spoofer maintains approximately fixed whatever initial carrier phase offset arises between its counterfeit signals and the authentic signals even as it shifts the code phase of its counterfeit signals to induce a position or timing deviation in the target receiver. This ability to approximately lock the relative (counterfeit-to-authentic) carrier phase even while shifting the relative (counterfeit-to-authentic) code phase enables the spoofer to evade some spoofing detection strategies that are designed to watch for the rapid amplitude variations caused by interacting authentic and counterfeit phasors of comparable magnitude when the authentic and counterfeit $\dot{\theta}$ values differ. However, when operating in the frequency-lock mode, the spoofer is limited to a code phase pulloff rate that lies within the target receiver's code tracking loop bandwidth, which can be as low as 0.05 Hz for carrier-aided code tracking [39]; otherwise, the target receiver will lose code lock on the counterfeit signals and the attack will be unsuccessful.

The spoofer makes no attempt to align its signals' carrier phases to those of the authentic signals. Nonetheless, by virtue of its carrier phase tracking and phase-locked signal generation, the real-time spoofer achieves nearly perfect phase coherence with the authentic signals during initial alignment (before tracking loop pulloff is attempted). More precisely, the differential Doppler frequency between each counterfeit signal and its authentic counterpart, as seen by the target receiver, is less than 0.01 Hz, a small offset that arises due to a linear approximation of the carrier phase trajectory over the $\sim$ 5-ms latency interval. In the post-processing testbed, differential Doppler is insignificant.

Precise carrier phase alignment would allow for more potent spoofing attacks, as it would enable generation of anti-phase signals that, if properly amplitude-matched, could annihilate each authentic signal. The spoofer could then generate a secondary ensemble of spoofing signals, in addition to the first anti-phase signal ensemble, which would be free of the telltale phase and amplitude variations caused by interaction with the authentic signals. However, such carrier phase alignment may only be practically possible under controlled laboratory conditions, as it would require spoofer-to-target relative position knowledge to within a small fraction of the carrier wavelength, which is approximately 19 cm for GPS L1. Indeed, the practical difficulty of carrier phase alignment in the field is the premise of the spoofing defense in [36].

2) Navigation Data Bit Prediction: To initialize an attack with an induced position, velocity, and timing solution that is indistinguishable from the authentic solution, it is not enough for the spoofer to achieve code-phase alignment with the authentic signals, it must also align its simulated navigation data bits with those of the authentic signals. However, due to processing, geometrical, and cable delays, it is impossible for the real-time spoofer to read the value of the navigation data bits off the air and replay them accurately and without delay. Indeed, this impossibility is precisely what makes navigation message authentication effective for GPS signal authentication, as discussed in [28] and [29].

Rather than read the navigation data bits off the air for immediate replay, the real-time spoofer takes advantage of the near perfect predictability of the navigation data that modulate the GPS L1 C/A signals. Over the course of a 12.5-minute superframe, the spoofer collects the data bits corresponding to each tracked GPS satellite. Alternatively, the spoofer can obtain the 12.5-minute superframe for each satellite from its control computer, which has access to a network of software-defined receivers of the type described in [40] and [41] that continuously generate intact superframes. Thereafter, the real-time spoofer compensates for its $\sim$5-ms processing delay, and for geometrical and cable delays, by predicting the value of the navigation data stream slightly more than 5 ms in advance. In this way, the spoofer can achieve meter-level alignment between its signals and the authentic ones at the location of a target receiver.

3) Variable Output Attenuation: Before exiting the real-time spoofer, counterfeit signals pass through an attenuator with a 31.5-dB range whose attenuation value can be set dynamically by the spoofer's control computer. This enables the spoofer to finely adjust the so-called spoofer power advantage, or the ratio of the power of the counterfeit signal ensemble to the power of the authentic signal ensemble as seen by the target receiver.

In the post-processing testbed, spoofer power advantage is adjusted by the digital combiner, which multiplexes blocks of n_s spoofing and n_a authentic samples, where n_s and n_a are user-defined integers. By properly adjusting the ratio n_s/n_a, a user can approximately achieve any reasonable spoofer power advantage.

4) Noise Padding: The signal ensemble generated by the testbed's spoofer contains only a modest amount of noise.

In other words, the native noise floor of the output signal ensemble is low—much lower than the noise floor present at the output of a high-quality GPS receiver's low-noise amplifier (LNA). To appreciate the consequence of this low native noise floor, consider that if the spoofer is configured to generate only a single output GPS L1 C/A signal, corresponding to a single pseudo-random number (PRN) code, the native C/N_0 of the output signal exceeds 60 dB-Hz. Of course, when more simulated GPS signals are added to the ensemble, the C/N_0 associated with any one of the signals drops due to multiple-access interference.

A low native noise floor would not be a problem for the spoofer if it were always configured to match the power of each counterfeit signal to that of the corresponding authentic signal at the RF input to the target receiver. In this case, the noise floor observed by the target receiver is essentially determined by the LNA in the receiver's own front-end.

But in some cases it may be advantageous for the spoofer to significantly overpower the authentic signals; for example, to eliminate interaction with them. In these cases, if the spoofer is generating a small number of simulated signals, the C/N_0 values registered by the target receiver for each received GPS signal become unnaturally high, owing to the low native noise floor of the spoofer's output ensemble. When generating a large number of signals—approximately 13 or more—the signals' mutual interference is sufficient to establish an appropriate noise floor for any particular spoofed signal.

To prevent unnaturally high C/N_0 values, the spoofer can be configured to add a variable level of "noise padding"—broadband interference—to its own output ensemble. In this way, the spoofer can dictate a maximum C/N_0 value for each of its output signals even while transmitting at high power.

5) Arbitrary Navigation Message Generation: In its default mode, the spoofer attempts to exactly match the data it modulates onto its counterfeit signals with the true navigation data on the corresponding authentic signals. This data-bit matching fails only in three circumstances: (1) during the first 18 seconds after a 2-hour GPS time boundary, when the GPS satellites begin broadcasting new ephemeris parameters in frames 1-3; (2) during a 12.5-minute superframe in which one or more satellites begin broadcasting new almanac data in frames 4-5, which occurs roughly once per day for each satellite; and (3) when the GPS satellites change reserved bits, which they occasionally do for reasons related to military receiver security. Other than in these situations, the spoofer's data bit matching is exact.

In some situations it may be advantageous for the spoofer to modulate its counterfeit signals with arbitrary data instead of matching the true navigation data streams bit-for-bit. This may be desirable, for example, to support a data manipulation attack, as in [42]. The testbed's spoofing device is capable of generating such arbitrary modulating data. For stealth and convenience, it does impose some structure on the data: (1) it maintains the legacy GPS subframe, frame, and superframe data format, (2) it populates the Handover Word (HOW) and the Telemetry Word (TLM) to match the authentic signals, (3) it respects data bits that are fixed in the GPS interface specification, and (4) it ensures that the data streams satisfy standard GPS L1 C/A parity checking.

6) SCER Attack Capability: The spoofer is capable of executing a so-called security code estimation and replay (SCER) attack. This attack targets cryptographic spoofing defenses in which an unpredictable (to the spoofer) security code modulates the transmitted GPS signal, whether as a component of the navigation data stream (navigation message authentication) or as higher-rate modulation.

When configured for a SCER attack, the spoofer seeks to estimate as best it can each security code chip value of each GPS signal that it intends to spoof. Its estimate for any particular chip is no better than a random guess at the beginning of the chip but improves rapidly thereafter. For a signal with received carrier-to-noise of $C/N_0 = 54$ dB-Hz, which is the highest that can be expected from a standard single-element hemispherical-gain-pattern GNSS antenna [43], the spoofer's chip estimation error becomes negligible after only 8 μs of averaging [29]; for more modest C/N_0, a few tens of μs is sufficient. As the spoofer obtains an estimate of each successive security code chip, it immediately injects this estimate into its signal replica generator, which is primed with up-to-date spreading code and carrier replicas. Thus the spoofer can approximately replicate even security-enhanced GNSS signals.

B. Defender

Opposite the spoofer in the real-time and post-processing testbeds sits a software-defined GNSS receiver that plays the role of defender. All digital signal processing downstream of the defender's RF front end is implemented in software on a (possibly multi-core) general-purpose processor. A software-defined receiver is well-suited for the role of defender because it is flexible enough to support rapid implementation and testing of a wide range of proposed defense strategies.

The particular software-defined receiver incorporated in the testbed, called GRID, is the result of nearly a decade of collaboration between the University of Texas at Austin and Cornell University [40], [41], [44], [45]. It has been designed for single- or multi-core platforms and has been implemented on Intel x86, Texas Instruments, and ARM processors. For efficient processing, key features of the receiver are its bit-wise parallel correlation strategy [46]–[48], its parallel architecture for multi-core implementation [45], and its use of SIMD instructions. Individually, and in combination, these features enable efficient signal processing despite the receiver being implemented on a general-purpose processor. On a 6-core processor, for example, GRID is capable of tracking 1150 parallel 5.7 Msps-sampled (real) GPS L1 C/A signals in real time.

Other key features of GRID useful for evaluating candidate signal authentication strategies are

1) access to raw 12-to-16 bit quantized digital samples prior to automatic gain control, which makes it possible to accurately measure changes in received in-band power;
2) a multi-tap correlation architecture, which allows examination of the correlation profile at arbitrary tap locations and with arbitrary density; and

3) sample-wise access to the product of the incoming signal and the local signal replica, which allows formulation of the detection statistic required in the SCER attack defense.

Details of these features will be introduced as needed in subsequent sections.

III. Evaluation of the Received Power Defense

For an important class of spoofing attacks, an admixture of authentic and spoofed GNSS RF signals is incident on the defender's antenna, which increases the total received power P_T in a GNSS band of interest beyond levels typically measured in the absence of spoofing. This observation suggests a low-complexity signal authentication strategy in which the defender chooses the null hypothesis H_0 (no spoofing attack underway) when P_T is within a nominal range, and the alternative hypothesis H_1 (spoofing attack underway) when P_T falls outside the nominal range. Indeed, this defense is proposed in [15] as "an extremely powerful means to detect spoofing, making spoofing no more of a threat than the much less sophisticated radio interference/jamming." This section evaluates the received power defense to determine whether it is indeed as potent as advertised.

A. Underlying Assumptions

Signal authentication based on P_T depends crucially on two assumptions, discussed below.

1) The Admixture Assumption: The received power defense assumes that a full admixture of counterfeit and authentic GNSS signals is present in the received band. If instead the attacker is able to partially or completely eliminate the authentic signals received by the defender, whether by annihilating these with anti-phase spoofing signals or, more simply, by covering the target antenna with an RF shield, then the attacker can prevent the defender's P_T from changing significantly during an attack.

The admixture assumption is reasonable in cases where (1) physical security prevents the attacker from gaining physical access to the defender's antenna, and (2) the attacker does not know the location of the defender's antenna to centimeter-level accuracy and so cannot mount an authentic-signal-annihilation attack. It is worth noting that some GNSS applications of practical interest violate these conditions: physical security obviously cannot be ensured when the attacker is in possession of the target receiver, as with a GPS ankle monitor or a vessel monitoring system [6], and the usual practice of mounting a GNSS antenna with open-sky access may enable an attacker to estimate its precise location, especially in the case of a static antenna.

Ref. [15] argues that, with proper calibration, "it should be possible to detect if the receiver is operating in open sky conditions or is blocked." But this is not the case, as one can appreciate with a simple thought experiment. Recall that the testbed's spoofer can adjust its output power over a 31.5 dB range in increments of 0.5 dB, and can artificially adjust the noise floor of its output signal ensemble. Moreover, the spoofer can independently measure the contribution to P_T due to ambient RF signals and background temperature and can accurately measure the relative C/N_0 of available GNSS signals. It follows that the spoofer can match both the absolute power of the authentic signal ensemble and the absolute C/N_0 value of each received GNSS signal. Thus, an attacker with physical access to a target receiver's antenna could slip a metal enclosure with an interior transmit antenna over the target antenna without causing significant variation in the defender's measured P_T and C/N_0 values. Incidentally, this "tin bucket" attack is also problematic for the pincer defense introduced in [36] and for defenses based solely on C/N_0 monitoring, as in [14].

2) The Small Unpredictable Variations Assumption: The received power defense also assumes that unpredictable variations in P_T, owing, for example, to solar radiation or to man-made but non-spoofing RF signal interference, are either small compared to the variations caused by spoofing, or rare. Otherwise, the false alarm rate for the spoofing detection test will be unacceptably high. This assumption is tested in [15] by monitoring variations in the automatic gain control (AGC) voltage, a proxy for $1/P_T$, over several days in quiescent (non-spoofing) conditions, and by comparing these with variations in AGC voltage observed during a live spoofing attack. In all cases tested, the AGC values during the spoofing attack stand out clearly against the quiescent AGC values whenever the target receiver's navigation solution is significantly affected. However, the attack executed in [15] does not permit determination of the minimum increase in P_T for a successful spoofing attack because the target receiver is always moving toward or away from the spoofer, so the spoofer cannot attempt a slow-pulloff low-transmit power attack. Moreover, [15] does not attempt to characterize common but unpredictable variations in P_T introduced by non-spoofing phenomena.

B. Detection Test

Signal authentication based on received signal power amounts to a binary hypothesis test in which the measurement P_T can be modeled as

$$H_0 : P_T = P_A + P_I + P_N, \tag{2a}$$

$$H_1 : P_T = P_C + P_I + P_N \tag{2b}$$

where $P_A = \sum_i P_{A,i}$ is the received signal power from an ensemble of n authentic GNSS signals in the absence of spoofing, $P_{A,i}$ being the power of the ith authentic signal; P_I is the received power from all man-made non-spoofing RF interference sources; $P_N = N_0 B$ is the received power from spectrally-flat receiver noise with density N_0 passing through a one-sided RF front-end bandwidth B; and P_C is the combined received power of the authentic and spoofing signals. The density N_0 is primarily determined by the noise figure of the receiver's first-stage LNA but also includes broadband noise due to solar and black-body radiation.

1) Effect of Coherence: Because of possible coherence between the received counterfeit and authentic signals, the combined signal power P_C is not simply a sum of the authentic and counterfeit signal powers. Let the total spoofing signal power that would be received in the absence of authentic

signals be $P_S = \sum_i P_{S,i}$, with $P_{S,i}$ being the spoofing signal power corresponding to the ith authentic signal. Further, let each $P_{S,i}$ be decomposed as $P_{S,i} = P_{Sc,i} + P_{Sn,i}$, where $P_{Sc,i}$ is the component of spoofing power that is coherent with the ith authentic signal and $P_{Sn,i}$ is the non-coherent component. The coherent component is assumed to have phase offset φ_i with respect to the ith authentic signal. One can now write P_C as

$$P_C = \sum_{i=1}^{n} \left[\sqrt{P_{A,i}} + \cos(\varphi_i)\sqrt{P_{Sc,i}}\right]^2 + \sin^2(\varphi_i) P_{Sc,i} + P_{Sn,i} \tag{3}$$

This expression indicates that, for each i, the noncoherent component $P_{Sn,i}$ adds directly to P_C, as does $\sin^2(\varphi_i)P_{Sc,i}$, which is the power in the coherent component that lies in phase quadrature to the authentic signal. By contrast, $\cos^2(\varphi_i)P_{Sc,i}$, which is the spoofing power component that is phase aligned with the authentic signal, does not add directly to P_C but instead interacts with the authentic signal as shown. For $k \in \mathbb{Z}$, the ith spoofing signal contributes maximally to P_C when $\varphi_i = k2\pi$ (phase alignment), minimally when $\varphi_i = (1+2k)\pi$ (anti-phase alignment), and power-additively—as if a purely noncoherent signal—when $\varphi_i = (1/2 + k)\pi$ (orthogonal alignment).

It is interesting to note that if the phase offsets φ_i are treated as independent random variables uniformly distributed on $[0, 2\pi]$, then the expected value of P_C is equivalent to the P_C that arises in the case of purely noncoherent spoofing signals; i.e.,

$$E[P_C] = P_A + P_S$$

Moreover, because the variance of P_C goes inversely with the number of signals n, it follows that for large n and $\varphi_i \sim U[0, 2\pi]$, P_C can be approximated as

$$P_C = P_A + P_S \tag{4}$$

However, the independence condition on the φ_i can be violated in practice by a spoofer with wavelength-level knowledge of the defender's antenna position, because in this case the spoofer can generate an ensemble of counterfeit signals at least some of whose φ_i will be similar. This has been demonstrated in the laboratory with this paper's testbed, as shown in Fig. 3. Outside the laboratory, however, violating (4) is only slightly less challenging for the spoofer than nulling the authentic signals.

2) Spoofing Power Advantage: For convenience, define

$$\eta \triangleq P_S / P_A \tag{5}$$

as the spoofing power advantage. Then P_C becomes a function of η, with $P_C(\eta = 0) = P_A$, and P_T can be rewritten as

$$P_T = P_C(\eta) + P_I + P_N \tag{6}$$

which, under the assumptions behind (4), becomes

$$P_T = (1 + \eta)P_A + P_I + P_N \tag{7}$$

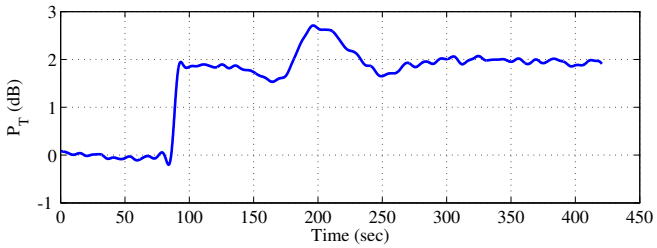

Fig. 3. Received power in a 2-MHz band centered at the GPS L1 frequency showing the onset of a spoofing attack using this paper's testbed, normalized by the average value of P_T prior to the attack. The attack begins with a sudden increase in P_T just before 100 seconds. Thereafter, the authentic power P_A and spoofing power P_S were maintained constant; thus, the oscillations in P_T can only be due to strong coherence between the spoofing and authentic signals with similar values of φ_i.

The hypotheses can now be written

$$H_0 : \eta = 0, \tag{8a}$$
$$H_1 : \eta \geq \eta_m \tag{8b}$$

where $\eta_m \geq 0$ is the minimum power advantage applied by a spoofer in an attack.

3) Simplifying the Composite Test: In view of (3), (6), and (8), deciding between H_0 and H_1 amounts to a composite hypothesis test in which the parameters η and $\varphi_i, i = 1, \ldots, n$ are simple under H_0 but can take on a range of values under H_1. The test can be reduced to a simple (non-composite) hypothesis test in two steps. First, since this paper's interest is in evaluating the strongest embodiment of the received power defense, let it be assumed that the defender knows the exact value of η. Second, assume the attacker does not have wavelength-level knowledge of the defender's antenna position, in which case it is reasonable to model the offsets φ_i as independent random variables uniformly distributed on $[0, 2\pi]$. Stacking these as $\boldsymbol{\varphi} = [\varphi_1, \varphi_2, ..., \varphi_n]^T$ and denoting the distribution of P_T under H_j by $p_{P_T|H_j,\boldsymbol{\varphi}}(\xi|H_j, \boldsymbol{\theta})$, $j = 0, 1$, one can integrate out $\boldsymbol{\varphi}$-dependence by

$$p_{P_T|H_j}(\xi|H_j) = \frac{1}{(2\pi)^n} \int p_{P_T|H_j,\boldsymbol{\varphi}}(\xi|H_j, \boldsymbol{\phi})\, d\boldsymbol{\phi} \quad j = 1, 2$$

where the multi-dimensional integral is taken over the range of $\boldsymbol{\varphi}$. The likelihood ratio can now be formed as

$$\Lambda \triangleq \frac{p_{P_T|H_1}(\xi|H_1)}{p_{P_T|H_0}(\xi|H_0)}$$

The optimal detection test compares Λ against a threshold [49]:

$$\Lambda \underset{H_0}{\overset{H_1}{\gtrless}} \tilde{\gamma} \tag{9}$$

This notation is interpreted as "choose H_1 if Λ exceeds $\tilde{\gamma}$; otherwise choose H_0." If the distribution of Λ is denoted $p_{\Lambda|H_j}(\lambda|H_j)$, $j = 0, 1$, then, for a chosen false alarm probability P_F, one sets $\tilde{\gamma}$ to satisfy

$$P_F = \int_{\tilde{\gamma}}^{\infty} p_{\Lambda|H_0}(\lambda|H_0)\, d\lambda \tag{10}$$

The resulting detection probability is

$$P_D = \int_{\tilde{\gamma}}^{\infty} p_{\Lambda|H_1}(\lambda|H_1)\ d\lambda \tag{11}$$

In many cases the test in (9) can be reduced to a simpler, equivalent test, e.g., by taking the log of both sides. Whatever quantity is ultimately compared against the final threshold, denoted γ, is called the detection statistic. For the special case where P_C, P_I, and P_N are modeled as Gaussian distributed, the problem becomes a simple location test in which the detection statistic reduces to P_T, which is itself Gaussian distributed [50]. Moreover, for small variations in P_T, the transformation to dB units via $P_T\ (\mathrm{dBW}) = 10\log_{10}(P_T)$ is approximately linear. Hence, for P_C, P_I, and P_N Gaussian, $P_T\ (\mathrm{dBW})$ can also be modeled as Gaussian.

C. Minimum Spoofing Power Advantage

Performance of signal authentication based on P_T depends crucially on η, P_A, P_I, and P_N, with the detection test becoming more powerful as η increases or as the variance in P_A, P_I, and P_N decreases. This section seeks to define η_m, a lower bound on η; the following section will examine P_A, P_I, and P_N.

1) Signal Model: By way of relating the parameters in (3), (5), and (6) to a signal model, consider an attack in which the received spoofing power is entirely coherent so that $P_S = \sum_i P_{Sc,i}$. Note that this implies the spoofer's output consists only of clean signal replicas with no quantization noise or noise padding. The defender's received signal at sampling instant t can then be represented by a complex baseband model as

$$\begin{aligned} r(t) = \sum_i \{&D_i(t)C_i\left[t-\tau_{ai}(t)\right]\exp\left[j\theta_{ai}(t)\right] \\ &+ \sqrt{\eta}D_i(t)C_i[t-\tau_{si}(t)]\exp\left[j\theta_{si}(t)\right]\} + I(t) + n(t) \end{aligned} \tag{12}$$

where, for the ith authentic signal, which is tracked in the receiver's ith channel, $D_i(t)$ is the navigation data, $C_i(t)$ is the spreading code, $\tau_{ai}(t)$ is the authentic signal's code phase, $\theta_{ai}(t)$ is the authentic signal's carrier phase, $\tau_{si}(t)$ is the spoofing signal's code phase, $\theta_{si}(t)$ is spoofing signal's carrier phase, $I(t)$ is a zero-mean complex process that models non-spoofing interference associated with P_I, and $n(t)$ is a zero-mean complex white Gaussian noise process that models the noise associated with N_0. This model remains a useful approximation even when mild quantization effects are present in the spoofing signals; it will be assumed to hold in the following analysis.

2) Successful Capture: A spoofer seeking to capture the defender's code and carrier tracking loops on each tracking channel while minimizing the likelihood of detection will operate with η near unity. Suppose that $\eta = 1$ and that, for the ith signal, $\tau_{si}(t) = \tau_{ai}(t)$ and $\theta_{si}(t) = \theta_{ai}(t) + \varphi_i(t)$. In this case, the received counterfeit and authentic GNSS signals are matched in amplitude and structure, differing only in carrier phase offset. If the spoofer now attempts pulloff of the defender's code phase tracking points in the default mode where code and carrier phase rates are related by (1), and if the spoofer maintains its carrier phase pulloff rate $\dot{\varphi}_i$ well below the defender's carrier phase tracking loop bandwidth B_L, then symmetry dictates that the spoofer's probability of successfully capturing the ith channel's code and carrier tracking loops is $p_{ci} = 0.5$.

In the absence of interference and noise [$I(t) = n(t) = 0$], $\eta > 1$ would be sufficient to guarantee capture of every channel's loops provided $|\dot{\varphi}_i/2\pi| \ll B_L, i = 1,\ldots,n$. But in the presence of interference and noise, $\eta > 1$ cannot guarantee capture even in the limit as $\dot{\varphi}_i \to 0$. This is because during pulloff there will be intervals during which $\varphi_i(t) \approx (1+2k)\pi, k \in \mathbb{Z}$, so that the counterfeit and authentic phasors will nearly annihilate each other. This phenomenon, which is redolent of severe ionospheric scintillation [51], can result in frequency unlock of the defender's carrier tracking loop, which for this paper's purposes is considered a failed capture.

For the ith signal, and for $\eta > 1$, the carrier-to-noise ratio during anti-phase alignment of counterfeit and authentic signals is

$$\frac{P_{A,i}(\eta-1)}{N_0}$$

To prevent frequency unlock, η must be chosen such that $P_{A,i}(\eta-1)/N_0 > \beta$, where β is the threshold value of C/N_0 required for frequency-unlock-free carrier tracking. This implies that, for all i, η must satisfy

$$\eta > 10\log_{10}\left[10^{(\beta - P_{A,i}/N_0)/10} + 1\right]\ \mathrm{dB}$$

in which η, β, and $P_{A,i}/N_0$ are expressed in dB. For a standard second- or third-order Costas-type GNSS carrier tracking loop with an update interval of 20 ms and $B_L = 5$ Hz, phase unlock begins below approximately $C/N_0 = 24$ dB-Hz [39], so one may take $\beta \approx 24$ dB-Hz as a conservative approximation for the frequency unlock threshold (the frequency unlock threshold is always below the phase unlock threshold). Thus, for a weak GNSS signal with $P_{A,i}/N_0 > 35$ dB-Hz, $\eta \geq \eta_u = 1.08$ (0.33 dB) would be required to prevent unlock.

3) Numerical Simulation and Testbed Experimentation: If $\eta \geq \eta_u$, then averaging within the tracking loops will ensure $p_{ci} \to 1$ as $\dot{\varphi}_i \to 0$. But a pulloff rate of zero is hardly useful for the spoofer. Within the more interesting interval $0 < |\dot{\varphi}_i/2\pi| \ll B_L$, the relationship between p_{ci}, η, and $\dot{\varphi}_i$ cannot be determined by a simple limiting case analysis. Moreover,a more comprehensive analytical examination of the code and carrier tracking loops is complicated by their stochastic, discrete, and nonlinear nature and by the counterfeit and authentic signal interaction. On the other hand, the closed-loop tracking behavior can be readily analyzed via Monte-Carlo simulation. Such a simulation has been carried out and has confirmed the general trends one might have expected: (1) for a fixed $|\dot{\varphi}_i/2\pi| \ll B_L$, p_{ci} quickly approaches unity as η increases beyond η_u, and (2) increasing η allows the spoofer to increase $|\dot{\varphi}_i/2\pi| \ll B_L$ while maintaining a fixed p_{ci}.

Apart from numerical simulation, the minimum value of η required for reliable capture has been determined experimentally via the testbed. On 34 independent trial attacks, each with

$n \geq 8$ authentic signals, it was found that $p_{ci} = 1$ whenever $\eta > 1.1$ (0.41 dB), provided $|\dot{\varphi}_i/2\pi| \ll B_L$ [52].

For purposes of this paper, it will be assumed that the spoofer always operates with $\eta \geq \eta_m = 1$ and that $\eta = \eta_m$ is enough to reliably capture the defender's tracking loops. This assumption is conservative as regards reliable capture because, as discussed above, capture only becomes reliable for $\eta \gtrsim 1.1$; yet it is optimistic as regards preventing adverse effects because a spoofer can cause a target receiver to output hazardously misleading data even when η is slightly less than unity. Nonetheless, for this paper it will be assumed that the attacker is not interested in uncontrolled adverse effects but in reliable capture requiring $\eta \geq \eta_m = 1$.

4) An Illustrative Scenario: It is instructive to roughly approximate the amount by which P_T changes between H_0 and H_1 given $\eta = 1$. Recall that P_T actually becomes the optimal detection statistic only when P_A, P_I, and P_N are modeled as Gaussian random variables, but in any case P_T closely approximates the optimal statistic. It follows that the detection test is powerful only if the increase in P_T from H_0 to H_1 is large compared to its random deviations under H_0 and H_1.

A typical outside-the-laboratory spoofing attack in which the assumptions behind (7) hold will yield the ratio

$$\frac{P_{T,1}}{P_{T,0}} = \frac{P_A(1+\eta) + P_I + P_N}{P_A + P_I + P_N} \tag{13}$$

of P_T under the two hypotheses. Consider an optimistic (for the defender) scenario in which $N_0 = -204$ dBW/Hz (a moderately low noise floor), $P_I = 0$ (no non-spoofing interference), $B = 2$ MHz (a narrow receiver bandwidth), and $P_A = -146$ dBW (consistent with an ensemble of typical-strength authentic GPS L1 C/A signals received in a $B = 2$ MHz band [9]). Despite the advantages to the defender in this scenario, $P_{T,1}/P_{T,0}$ is only 0.93 dB when $\eta = 1$. For $N_0 = -201$ dBW/Hz, which is more realistic for a commercial-grade GNSS receiver, $P_{T,1}/P_{T,0}$ falls to 0.56 dB. Roughly speaking, then, powerful received-power-based signal authentication requires that random fluctuations in P_T be substantially smaller than 1 dB. This is a restatement of the small unpredictable variations assumption.

D. Characterization of P_A, P_I, and P_N

The causes of variations are different for each of P_A, P_I, and P_N. Some variations can be accurately predicted by the defender, and so can be treated as deterministic, whereas others are not practically predictable and must be modeled as random. An analytical treatment of these random variations is not possible, as they are highly device-, site-, and time-specific. Therefore, this section appeals to empirical study.

Fig. 4 shows the RF spectrum centered at the GPS L1 frequency as seen by a high-quality static antenna and wide-bandwidth RF front end combination. The power spectral density is estimated by generating periodograms using Welch's method on 100-ms intervals of raw complex samples and then averaging over 100 of these. The characteristic peak resulting from the noncoherent combination of approximately 12 GPS L1 C/A signals is visible above the noise floor. Two bands are shown centered at L1, a 2-MHz band, which contains 90% of the L1 C/A signal power, and a 10-MHz band, which contains 98%. No spurious signals are visible in either band, which implies that $P_I \approx 0$.

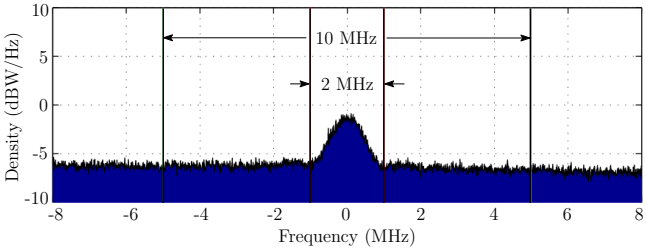

Fig. 4. Power spectrum centered at the GPS L1 frequency as estimated from data captured via a high-quality static antenna and RF front end combination in a moderately quiet outdoor RF environment on the rooftop of the WRW building on the UT Austin campus. Bands for 2- and 10-MHz power measurements are shown. The power density scale has been centered just above the GPS L1 C/A peak for ease of viewing. In absolute units, the noise floor sits at approximately -204 dBW/Hz.

Summing the 100-ms periodograms over the bands indicated results in a time series of power measurements. Fig. 5 shows a two-day interval of P_T in the 2-MHz band, which reveals marked diurnal variations, the result of diurnal patterns in temperature, solar radiation, and the overhead satellite constellation. Even though the record's diurnal repeatability is evidently only good to roughly 0.3 dB, its predictability given knowledge of local temperature and satellite orbital ephemerides is better than this. Fig. 6 offers an expanded view of a 5-minute interval, showing both the 2- and 10-MHz traces. The different size of the variations in the two traces at time scales less than about 150 seconds indicates that these originate in P_A, not P_N. They are likely due to multipath effects at the carrier phase level caused by reflections off nearby surfaces and by atmospheric diffraction and refraction. Close examination of multi-day records such as those in Fig. 5 reveals that these variations do not repeat appreciably at the solar or sidereal day. Data from two other static sites were examined, with similar behavior noted. Thus, it appears that the practically unpredictable variations in P_T about L1 have root-mean-squared deviations of at least 0.1 dB for a 2-MHz band and 0.05 dB for a 10-MHz band.

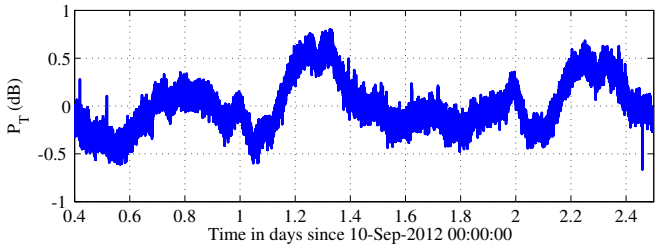

Fig. 5. A two-day record of received power in the 2-MHz band shown in Fig. 4, normalized by the average value of P_T over the interval.

Suppose that P_T, in dB units and with its mean under H_0

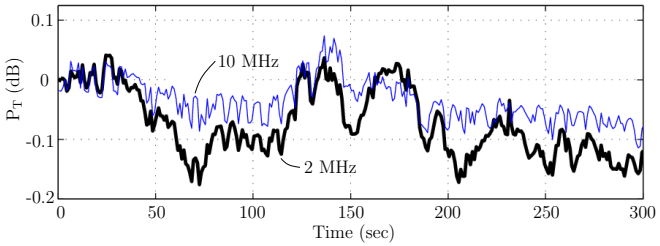

Fig. 6. A five-minute record of received power in the 2- and 10-MHz bands shown in Fig. 4, normalized by the initial values of P_T in each band.

removed, is taken as the detection statistic and modeled as

$$H_0 : P_T \text{ (dBW)} \sim \mathcal{N}(0, 0.1), \tag{14a}$$

$$H_1 : P_T \text{ (dBW)} \sim \mathcal{N}(0.56, 0.1), \tag{14b}$$

where the mean value under H_1, 0.56, is taken from the discussion of $P_{T,1}/P_{T,0}$ in Section III-C4. Choice of an acceptable P_F depends on the cost of a false alarm, which may range from a site visit to the grounding of an aircraft. As a reasonable value, assume only one false alarm per year is acceptable. Then if, due to the time correlation evident in Fig. 6, an independent test occurs every 150 seconds, a once-per-year alarm corresponds to $P_F = 4.75 \times 10^{-6}$. For this P_F, the decision threshold calculated via (10) is $\gamma = 0.44$ dBW, and the detection probability is $P_D = 0.88$.

This value of P_D gives reason to be optimistic about signal authentication based on P_T for static GNSS receivers. Such performance depends, however, on the distribution of P_T having exponentially decaying tails. In practice, there are at least two phenomena that can cause P_T to routinely take on values that would be exceedingly improbable under a Gaussian distribution: solar radio bursts and non-spoofing interference.

1) Solar Radio Bursts: Recall that P_N represents the contribution to P_T due to spectrally-flat receiver noise. It can be related to the receiver and antenna noise temperatures T_R and T_A (in degrees Kelvin) by

$$P_N = BN_0 = k_B B(T_R + T_A) \tag{15}$$

where k_B is Boltzmann's constant.

Unpredictable variations in T_R arise due to random fluctuations in noise sources internal to the receiver, primarily those in the first-stage LNA. These are small enough they do not contribute significantly to the ~ 0.1 dB variations in P_T noted previously for static antennas.

Variations in T_A arise due to antenna motion (as more or less warm earth radiation is visible), antenna blockage (e.g., an increase in T_A due snow accumulation [15]), and variable solar radiation. All these would be difficult or impossible for a stand-alone (non-networked) GNSS receiver to predict. Focus here will be on solar radiation as its effect is least site-specific: all GNSS receivers in the sunlit portion of the earth are similarly affected.

Solar radio bursts can cause large and sudden variations in P_N, as exemplified by the December 2006 storm, which led to 10-17 dB increases in P_N [53]. The relevant question as regards P_T-based GNSS signal authentication is how often a burst event would cause P_T to exceed the detection threshold, causing a false alarm. This question is answered in Table I for three different values of the threshold γ.

TABLE I
TIME BETWEEN THRESHOLD-EXCEEDING SOLAR RADIO BURST EVENTS FOR VARIOUS VALUES OF THE DETECTION LEVEL γ

	Threshold Value		T_e (days)	
γ (dB)	T_{As} (K)	S_1 (SFU)	Solar max.	All years
0.44	40.9	1560	9.2	22
0.93	91.3	3488	17.3	42.9
1.5	157.7	6022	26.5	67.4

Table I is interpreted as follows. Assume $P_I = 0$ and let $T_A = T_{A0} + T_{As}$, where T_{As} is the portion of T_A due to solar radiation. Each γ value can then be related to a threshold T_{As} by

$$\gamma \text{ (dB)} = 10 \log_{10} \left[\frac{P_A + k_B B(T_R + T_{A0} + T_{As})}{P_A + k_B B(T_R + T_{A0})} \right]$$

assuming the following reasonable parameter values: $P_A = -146$ dBW, $B = 2$ MHz, $T_R = 188$ K, $T_{A0} = 100$ K. Each T_{As}, in turn, is related to a threshold solar flux density S_1 by

$$S_1 \text{ (SFU)} = \frac{2k_B T_{As}}{A_e 10^{-22}}$$

where the effective antenna area is taken to be $A_e = 7.23 \times 10^{-3}$ m^2, which is a good approximation for a single-element GNSS antenna, and the additional factor of 2 in the numerator reflects the assumption that only half the total-polarization solar radiation contributes to T_{As} through a GNSS antenna, which is designed to received right-hand circularly polarized signals [54]. The factor 10^{-22} converts W/m^2/Hz to solar flux units (SFU). The resulting S_1 values listed in Table I are those above which a spoofing detector based on P_T would declare H_1 for the corresponding γ. As a final step, the model

$$N(S > S_1, \nu_1, \nu_2)$$

from [55] is invoked (with the correction factor C_{geo}) to approximate the total number of bursts exceeding S_1 in the frequency range $[\nu_1 = 1 \text{ GHz}, \nu_2 = 1.7 \text{ GHz}]$ over a 40-year historical period. This is used to estimate T_e, the time between triggering events, for solar maximum years and for all years.

Table I makes clear that solar radio bursts are problematic for signal authentication based solely on P_T. Under the model in (14), the threshold $\gamma = 0.44$ dB leads to a respectable $P_D = 0.88$ for a once-per-year false alarm. Accounting for solar radio bursts, the P_D remains approximately unchanged, but the false alarm rate rises to once every 9 days during solar maximum, or once every 22 days on average across the full solar cycle. This rate would be unacceptably high for many applications. Worse yet, there is little refuge in higher γ values as there would be for a P_T distribution having exponentially decaying tails. At $\gamma = 0.93$ dB, which would only yield $P_D = 0.5$ even under the higher-sensitivity spoofing attack scenario in Section III-C4, the false alarm rate is still greater than once every two months. Even for $\gamma = 1.5$ dB, which would offer no detection power against a spoofing attack with $\eta = 1$, and only $P_D = 0.5$ for $\eta = 1.7$, the long high-side tail of the true

P_T distribution prevents the false alarm rate from dropping to less than one event in three months.

If these false alarm rates are unacceptable, as they would be for many applications, then a spoofer could operate without fear of detection so long as it set η near unity. One may object to this conclusion by pointing out that spoofing alarms could be dismissed during known solar radio burst events, which can be independently monitored—even predicted (see http://www.swpc.noaa.gov/). But this offers little protection, for a clever attacker could time his attack to coincide with the arrival of a sizable burst.

2) Non-Spoofing Interference: Laying aside concerns due to solar radio bursts, one must also consider the effect of non-spoofing interference on P_T-based signal authentication. Such interference, whose received power is represented by P_I, ranges from unintentional in-band harmonics to intentional jamming [56]. It can affect both stationary and moving GNSS receivers, though the variance of P_I will generally be higher for moving receivers.

The mean and variance of P_I are context specific, but both tend to increase with population density [57]. In recent years, interference due to so-called personal privacy devices has become an increasing concern [58], [59]. Current use of these jammers along major highways results in P_T spikes that, for nearby receivers, would violate any of the thresholds considered in Table I. Moreover, as shown in Fig. 7, the jamming profiles seen at closely-spaced sites are different enough that there will remain a substantial unpredictable P_I component even if local monitoring is in place.

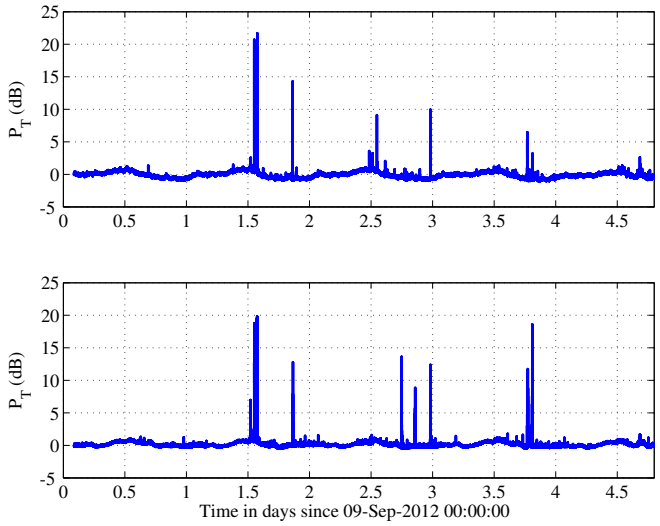

Fig. 7. Received power in the 10-MHz band centered at GPS L1 at two sites 1 km apart that straddle State Highway 1, west of Austin, TX. Top panel: Data from site located at the Center for Space Research. Bottom panel: Data from site located at Applied Research Laboratories. Both traces are normalized by the average value of P_T over the interval.

One might argue that it is perfectly appropriate for a spoofing detector to alarm in the presence of an intentional jammer, but the consequences of spoofing can be much more malign than those of jamming, and so it behooves a defender to distinguish the two.

Note that non-spoofing interference is not only a problem for P_T-based signal authentication but for all GNSS signal authentication methods that depend on constraining the spoofer to low values of η, such as the pincer defense [36]. This defense is less sensitive to solar radio bursts than P_T-based signal authentication, but equally likely to declare a false alarm in the face of strong non-spoofing interference.

E. Evaluation Summary

Even granting the full signal admixture assumption, it appears that, contrary to the claim made in [15], spoofing detection based solely on received power P_T is inadequate for GNSS signal authentication, for two reasons: (1) the increase in P_T due to spoofing can be small (less than 1 dB), and (2) a long tail in the distribution of P_N due to solar radiation causes high P_F for any reasonable P_D and, for receivers in urban areas, the same can be true for P_I due to non-spoofing interference. These conditions amount to a violation of the small unpredictable variations assumption.

Despite its weakness, a P_T-based defense remains a useful component of GNSS signal authentication, as it prevents an attacker from employing an arbitrary η. It is best thought of as a necessary, but not sufficient, test for GNSS signal authentication. For increased potency, P_T testing can be combined with a correlation distortion test, as in [36], a cryptographic test, as in [28], [29], or another substantially independent and complementary test. Note that jointly testing for unusual P_T and C/N_0 values is only slightly better than testing P_T alone: at the expense of a slightly higher η, a spoofer can inject noise padding to ensure that its signals' C/N_0 values match those of the authentic signals.

IV. Evaluation of the SCER Attack Defense

The SCER attack defense, originally developed in [29], assumes that the authentic broadcast GNSS signals have been modulated with a signal-specific binary security code that is unpredictable to the spoofer but verifiable by the defender (possibly after a delay). Unable to predict the security code, the spoofer resorts to modulating its counterfeit signal replicas with security code chips estimated on-the-fly. The key to defending against a SCER attack is a detection statistic sensitive to the high error variance of the spoofer's security code chip estimates in the moments immediately following each unpredictable chip transition. Ref. [29] develops such a statistic, describes its distribution under H_0 and H_1, and offers preliminary results using this paper's testbed. This section explains how the detection statistic is generated in practice within a GNSS receiver and offers a more extensive empirical evaluation of the SCER attack defense.

A. Detection Test

A single-signal SCER attack can be modeled by the following hypothesis pair for the samples Y_k output by the defender's RF front end during the interval spanned by the lth security code chip:

$$H_0 : Y_k = W_l s_k + N_k, \tag{16a}$$

$$H_1 : Y_k = g\left[\alpha \hat{W}_l(n_{lk}) s_k + N_k\right] \tag{16b}$$

Under hypothesis H_0, the received signal is an authentic GNSS signal with security code chip value W_l and underlying signal $s_k = c_k \cos(2\pi f_{\rm IF} t_k + \theta_k)$, where c_k is the signal's binary spreading code, $f_{\rm IF}$ is the intermediate frequency in Hz, and θ_k is the beat carrier phase. The noise samples N_k are modeled as independent and Gaussian. Under hypothesis H_1, the received signal is a spoofer-generated exact counterfeit of s_k modulated by an estimate $\hat{W}_l(n_{lk})$ of the lth security code chip. The index n_{lk} represents the number of samples that contribute to the spoofer's estimate of W_l. The coefficient α is the spoofing amplitude factor, which is proportional to $\sqrt{\eta}$, and g is the automatic gain control factor imposed by the RF front end to maintain constant power in Y_k.

Ref. [29] offers further details on the model in (16) and formulates a detection statistic appropriate for defending against a SCER attack. The current paper illustrates how this statistic is generated within a GNSS receiver. For clarity of presentation, assume the security code is carried in the navigation data stream so that each unpredictable security code chip W_l is also a navigation data symbol. In other words, assume a navigation message authentication security scheme [28]. Further assume that the receiver's accumulation (pre-detection) interval is equivalent to the length of W_l. Then the detection statistic L can be generated as shown in Fig. 8.

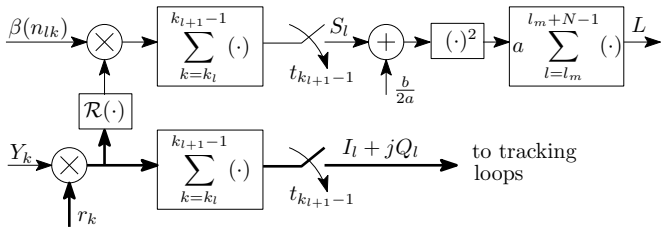

Fig. 8. Block diagram illustrating how generation of the SCER attack statistic L relates to standard GNSS signal correlation. Thick lines denote complex signals, whereas thin lines denote real-valued signals.

By way of further explanation, consider the two signal paths shown in Fig. 8. The lower path is the standard matched-filter-type correlation operation commonly implemented in GNSS receivers. The product of the incoming samples Y_k and a complex local signal replica $r_k = W_l \hat{c}_k \exp[-j(2\pi f_{\rm IF} t_k + \hat{\theta}_k)]$ is accumulated over the interval spanned by W_l to produce the prompt complex correlation products $I_l + jQ_l$ that get fed to code and carrier tracking loops. The code tracking loop also ingests correlation products from identical paths—not shown—having early and late versions of $\hat{c}_k$.

The upper path in Fig. 8 produces the SCER attack detection statistic L. The real part of the product $Y_k r_k$ is multiplied by a smooth weighting function $\beta(n_{lk})$, defined in [29], that gives full weight to the k_lth sample but decays rapidly toward zero for subsequent samples. This weighting has the effect of suppressing those samples over which the error variance in the spoofer's security code chip estimate $\hat{W}_l$ has become small because the spoofer has had sufficient time to obtain an accurate estimate of W_l; only the early high-variance samples are useful in distinguishing H_1 from H_0. The weighted product $\beta(n_{lk})\mathcal{R}(Y_k r_k)$ is accumulated over the interval spanned by W_l to produce the single-chip detection statistic S_l, N of which are biased, squared, and accumulated as shown to produce the final statistic L. The constants a and b are related to the theoretical mean μ_j and variance σ_j^2 of S_l under $H_j,\ j = 0, 1$ by

$$a = \frac{1}{\sigma_0^2} - \frac{1}{\sigma_1^2}, \qquad b = 2\left(\frac{\mu_1}{\sigma_1^2} - \frac{\mu_0}{\sigma_0^2}\right)$$

B. Test Setup

Due to its ~ 5-ms processing latency, the real-time spoofer, in its current form, is not capable of a near-zero-latency SCER attack in which the spoofer's output security code chip estimates are approximately aligned with those of the authentic security-code-enhanced signals when received by the defender. Note that although a zero-latency attack is physically impossible for a real-time system, a near-zero-latency attack (e.g., less than 50 ns latency) could be achieved in real time with an FPGA-based real-time spoofer. For the SCER attack results presented in this paper, the post-processing testbed's digital I/O spoofer was used, which can be configured to mount a SCER attack with arbitrary latency. To permit evaluation of the most-potent limiting case, the digital I/O spoofer was configured to mount a zero-latency attack.

The attack proceeded as follows. The digital I/O spoofer ingested authentic recorded GPS L1 C/A data and, treating the ± 1-valued 20-ms navigation data bits as if they were unpredictable security code chips, generated a maximum *a posteriori* (hard-decision) estimate for each chip. Near the beginning of each chip, when the spoofer had few signal samples on which to base its estimate, these chip values would switch wildly between -1 and 1. But with each successive sample received, the error variance of the spoofer's chip estimate would diminish until, after about 100 μs, the estimate would become virtually certain. The spoofer continuously modulated each of 8 constituent spoofing signals in its output ensemble with the corresponding chip estimate trains.

The spoofer began its attack with its counterfeit signals approximately code-phase-aligned and data-aligned to the authentic signals. After maintaining this alignment for several hundred seconds, it attempted pulloff of the defender's tracking loops, stopping once it had attained an offset of 175 μs with respect to the authentic signals. Due to the orthogonality of the GPS C/A codes, there was no significant interplay between the authentic and counterfeit signals at this offset.

The digital I/O spoofer's output data were sample-wise multiplexed with the original authentic data to produce a digital data stream containing the composite spoofing and authentic signal ensembles. The multiplexing ratio was adjusted so that $\eta \approx 1.2$. A preliminary segment of the data was left free of spoofing to allow testing of the defender's ability to detect the onset of attack.

The combined data stream was routed to the testbed's digital-input software-defined receiver, acting as defender, which tracked the signals present and produced samples equivalent to the product $Y_k r_k$ in Fig. 8. The real parts of these samples were weighted by an appropriate $\beta(n)$ and accumulated to generate a sequence of chip-level statistics S_l. Batches of $N = 400$ S_l were combined to produce a full

detection statistic L every 8 seconds during the course of the experiment.

All signals tracked by the spoofer had spoofer-measured carrier-to-noise ratios $(C/N_0)_s \geq 46$ dB-Hz whereas, due to the way multiplexing was effected, the authentic signals tracked by the defender prior to attack had defender-measured carrier-to-noise ratios $40 < (C/N_0)_r < 42$ dB-Hz. Thus, the spoofer enjoyed at least a 4 dB carrier-to-noise advantage over the defender in the attack, which, for the defender, represents a challenging attack scenario. In the formulation of L, the defender's assumed values for $(C/N_0)_s$ and $(C/N_0)_r$, which influence μ_j and σ_j, $j = 1, 2$, and, by extension, the theoretical distributions $p_{L|H_j}(\xi|H_j)$, $j = 1, 2$, were taken to be approximately the true values of $(C/N_0)_s$ and $(C/N_0)_r$. The defender's assumed value of η was taken to be $\eta = \eta_m = 1$, not far from the true $\eta = 1.2$. Thus, the defender's model for the distribution of L, upon which its decision threshold for each signal was based, was approximately equal to the true distribution of L for that signal, except during the initial aligned stage of the attack over which interaction of the spoofing and authentic signals unavoidably violated the model in (16). The defender's detection threshold was set such that $P_F = 10^{-4}$.

C. Test Results

The following test results are expressed in terms of the empirical distribution of L at various stages of a SCER attack. Typical results will be presented first, followed by discussion of less typical results.

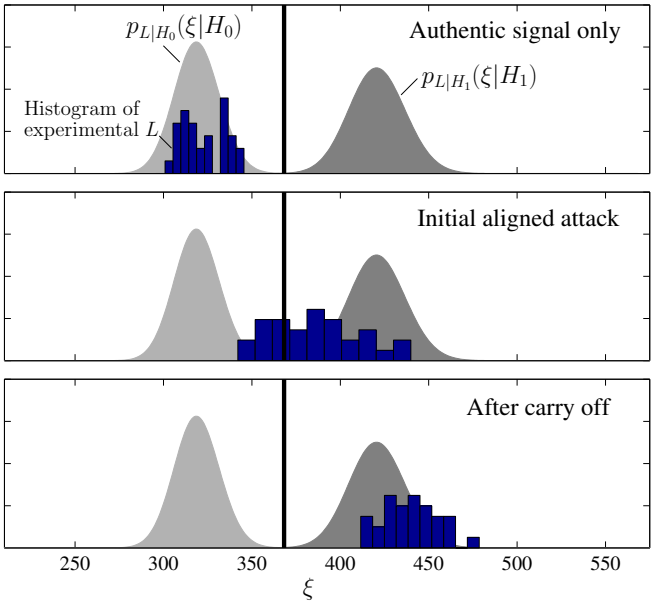

Fig. 9. Histograms of experimentally-generated detection statistics L (bar plots) compared with the detection threshold (thick vertical line) and the theoretical distributions $p_{L|H_j}(\xi|H_j)$, $j = 0, 1$ at various stages of a zero-delay SCER attack on the signal corresponding to PRN 17.

1) Typical Results: The top panel in Fig. 9 shows the attack prelude during which only the authentic signal was present. At this stage, the histogram of L values exhibits good correspondence with the theoretical null-hypothesis probability distribution $p_{L|H_0}(\xi|H_0)$. The center panel shows the situation during the initial stage of the attack when the authentic and spoofing signals were aligned to within a small fraction of the ~ 1-μs spreading code chip interval. Because the counterfeit and authentic signals in this test were so nearly matched in power, this stage saw strong interaction between them in the defender's complex-valued prompt correlator. Such interaction violates the either/or assumption of (16); nonetheless, the detection statistic exceeds the threshold more than half the time. However, instead of clustering within $p_{L|H_1}(\xi|H_1)$, the histogram exhibits spreading. Fig. 10 shows a time history of L during this stage of the attack. The slow changes in L are driven by variations in the relative carrier phase of the interacting authentic and spoofing signals.

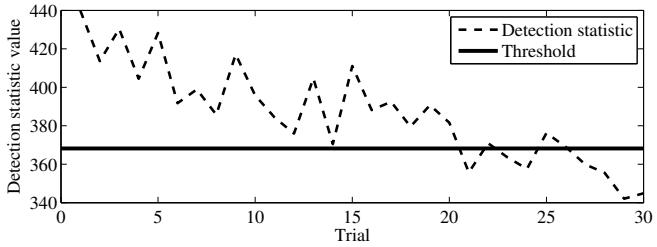

Fig. 10. A time history of the defender-measured value of the decision statistic L during the aligned stage of the attack on PRN 17. Each trial represents an 8-second interval.

After the spoofer has successfully carried off the defender's tracking points and the authentic and spoofed correlation peaks are separated by more than two spreading code chips, the model in (16) again becomes valid. The bottom panel of Fig. 9 shows that at this stage the detection statistic clearly clusters beyond the detection threshold and roughly within the theoretical $p_{L|H_1}(\xi|H_1)$ distribution. It should be noted that in the experiment the post-pulloff C/N_0 value measured by the defender did not change significantly relative to the measured C/N_0 prior to the attack. Thus, a naive spoofing detection strategy that triggers on changes in C/N_0 would have failed to detect this attack.

The favorable results shown in Fig. 9, together with those originally presented in [29], are fairly typical—they are representative of 2/3 of the results from similar experiments conducted on the testbed at various values of $(C/N_0)_r$ and $(C/N_0)_s$.

2) Atypical Results: Figs. 11 and 12 show results representative of the remaining 1/3 of the cases studied. As with the previous results, the empirical histograms of L under H_0 exhibit good agreement with the theoretical $p_{L|H_0}(\xi|H_0)$ (top panels). The histograms during the initial aligned attack (center panels) are to the left of the threshold [Fig. 11] or spread widely [Fig. 12], yet not atypical given the various ways that the counterfeit and authentic signals can interact at this stage. However, under H_1 (bottom panels), the empirical histograms are unusual: they are wider than the theoretical $p_{L|H_1}(\xi|H_1)$, and, in the case of Fig. 12, lower in mean value. This mismatch has the effect of reducing P_D to 0.87 for the case in Fig. 11 and to 0.46 for Fig. 12. The reason for this mismatch is unclear,

as there was no significant interaction between authentic and counterfeit signals at this stage of the attack.

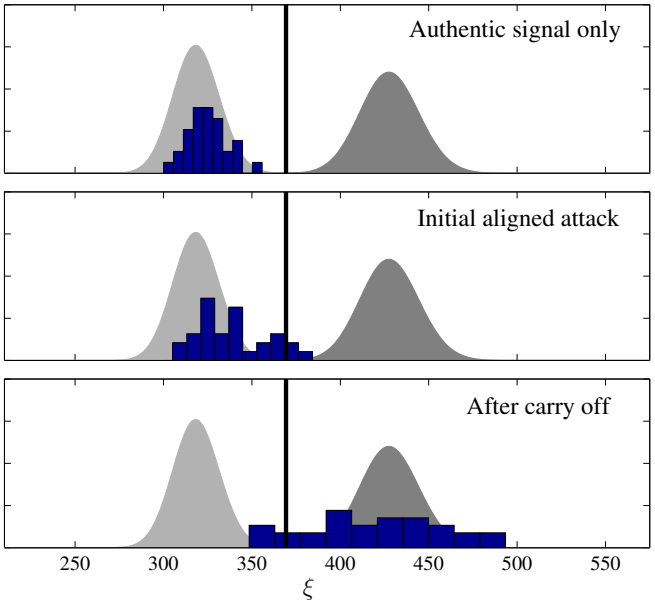

Fig. 11. As Fig. 9 except for PRN 27.

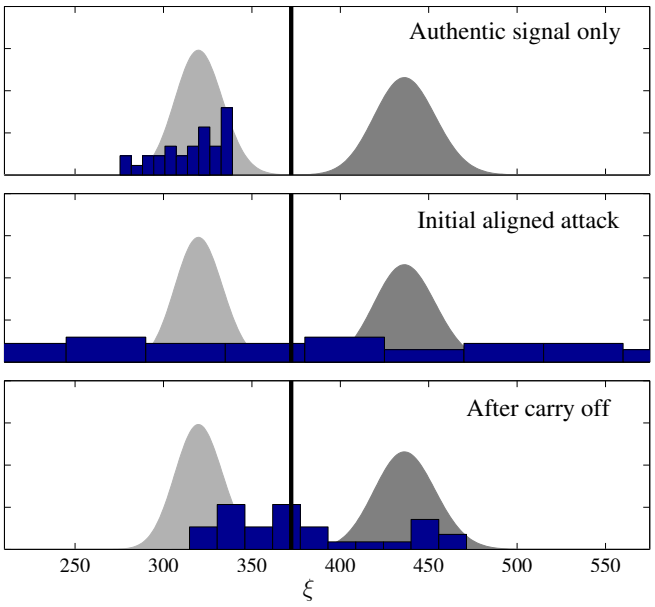

Fig. 12. As Fig. 9 except for PRN 4.

D. Evaluation Summary

Experimental results indicate close agreement between the empirical and theoretical distributions of L under H_0. This implies that the false alarm rate for the SCER attack defense is consistent with the value of P_F used to set the detection threshold. If the value $P_F = 0.0001$ chosen in the experiments is unacceptably high for a given application, P_F can be lowered while maintaining a useful P_D: for a low-rate security code, $P_F = 10^{-6}$ results in $P_D > 0.85$ [29]. Similarly, in 2/3 of cases studied there was close agreement between the empirical and theoretical distributions of L under H_1, which implies that the theoretical value of P_D, which was near unity for all the experimental scenarios studied, can be approximately reached in practice. Even in atypical cases of disagreement, P_D remained above 0.46. Thus, compared to the received power defense, the SCER attack defense is significantly more powerful.

Nonetheless, the SCER attack defense has three weaknesses. First, during the initial stage of a signal-aligned attack, L can remain below the detection threshold over an extended interval due to interaction between the authentic and counterfeit signals [cf. Fig. 11, center panel]. One might think that poor P_D is irrelevant at this stage given that the spoofer has not yet attempted pulloff, but it turns out that if a majority of signals are being spoofed the multipath-like effects of aligned counterfeit and authentic signal interaction can cause navigation errors of several tens of meters. Of course, in this case the likelihood that at least one channel's L rises above the detection threshold remains quite high, so one may consider this a minor weakness.

The second weakness of the SCER attack defense concerns the spoofing power advantage η. It is shown in [29] that a defender can maintain P_D above 0.9 even under a challenging SCER attack scenario so long as η is known. When the defender significantly underestimates η, however, P_D can fall precipitously for low $(C/N_0)_r$. The defender could address this weakness by estimating η via observation of P_T, using (13) and taking $\eta_m = 1$ as a lower bound on the estimate. This amounts to a generalized likelihood ratio test with η as the composite parameter to be estimated [49]. Note that, under this strategy, an increase in P_T due to a solar radio burst or non-spoofing interference would not significantly affect P_F.

The third and most significant weakness of the SCER attack defense is that it fails in the case of a near-zero-latency pure replay (meaconing) attack because in this case $\hat{W}_l = W_l$. While one should not expect a defense designed for SCER attacks to also detect a pure replay attack, it nonetheless remains true that a pure replay attack is easy to mount—much easier than a SCER attack—and, while not enjoying the same flexibility as a SCER attack to dictate an erroneous navigation and timing solution, is dangerously effective. To address this weakness, the SCER attack defense could be combined with the pincer defense [36], which is effective against a pure replay attack. However, like the received power defense, the pincer defense is prone to false alarms in the face of a large increase in P_T not related to spoofing.

V. Conclusions

An experimental testbed for developing and evaluating GNSS signal authentication techniques has been described and used to evaluate two candidate signal authentication techniques. It was shown that the first technique, the received power defense proposed in [15], fails to detect a spoofing attack when the spoofing power advantage $\eta \approx 1$ and when the false alarm probability $P_F < 10^{-6}$. Even when $P_F = 10^{-4}$, which would result in approximately one false alarm every

17 days during solar maximum, the detection probability P_D remains below 0.5. Nonetheless, the received power defense remains useful for detecting unsophisticated spoofers that resort to $\eta \gg 1$.

The SCER attack defense proposed in [29] was also evaluated, assuming a low-rate security code consistent with navigation message authentication. In most cases, the empirical P_F and P_D matched the modeled values, which ensured $P_D \approx 1$ for $P_F = 10^{-4}$ and $P_D > 0.85$ for $P_F = 10^{-6}$. However, in some cases the empirical P_D dropped below the theoretical P_D, sometimes as low as 0.5 for $P_F = 10^{-4}$. The SCER attack defense may also suffer from low P_D during the initial stage of an aligned attack, though if several signals are spoofed the chance of at least one channel alarming remains high. For good performance, the SCER attack defense should continuously estimate η from measurements of the received power P_T. The most significant weakness of the SCER attack defense is its inability to detect a pure replay (meaconing) attack, which, while not as flexible as a SCER attack, is nonetheless potent and dangerous. However, it should be noted that all cryptographic GNSS signal authentication schemes, even those based on high-rate military-style security codes, are vulnerable to pure replay attacks.

REFERENCES

[1] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. O'Hanlon, and P. M. Kintner, Jr., "Assessing the spoofing threat: Development of a portable GPS civilian spoofer," in *Proceedings of the ION GNSS Meeting*. Savannah, GA: Institute of Navigation, 2008.

[2] 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*, vol. 5, no. 3-4, pp. 146–153, 2012.

[3] 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*, 2014, http://dx.doi.org/10.1002/rob.21513.

[4] John A. Volpe National Transportation Systems Center, "Vulnerability assessment of the transportation infrastructure relying on the Global Positioning System," 2001.

[5] N. S. W. Center, "Global positioning system impact to critical civil infrastructure (GICCI)," Mission Assurance Division, Naval Surface Warfare Center, Tech. Rep., 2009.

[6] U. Kroener and F. Dimc, "Hardening of civilian GNSS trackers," in *Proceedings of the 3rd GNSS Vulnerabilities and Solutions Conference*. Krk Island, Croatia: Royal Institute of Navigation, Sept. 2010.

[7] Department of Homeland Security, "National risk estimate: Risks to U.S. critical infrastructure from Global Positioning System disruptions," November 2012, FOUO: No Public Version Available.

[8] U.S. Government Accountability Office, "Unmanned Aircraft Systems: Measuring progress and addressing potential privacy concerns would facilitate integration into the national airspace system," GAO-12-981, September 18, 2012, http://www.gao.gov/products/GAO-12-981.

[9] J. J. Spilker, Jr., *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996, ch. 3: GPS Signal Structure and Theoretical Performance, pp. 57–119.

[10] GPS Directorate, "Systems engineering and integration Interface Specification IS-GPS-200G," 2012, http://www.gps.gov/technical/icwg/.

[11] European Union, "European GNSS (Galileo) open service signal in space interface control document," 2010, http://ec.europa.eu/enterprise/policies/satnav/galileo/open-service/.

[12] B. M. Ledvina, W. J. Bencze, B. Galusha, and I. Miller, "An in-line anti-spoofing module for legacy civil GPS receivers," in *Proceedings of the ION Internation Technical Meeting*, San Diego, CA, Jan. 2010.

[13] K. D. Wesson, D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, "An evaluation of the vestigial signal defense for civil GPS anti-spoofing," in *Proceedings of the ION GNSS Meeting*, Portland, OR, 2011.

[14] V. Dehghanian, J. Nielsen, and G. Lachapelle, "GNSS spoofing detection based on receiver C/N_0 estimates," in *Proceedings of the ION GNSS Meeting*. Nashville, Tennessee: Institute of Navigation, 2012.

[15] D. M. Akos, "Who's afraid of the spoofer? GPS/GNSS spoofing detection via automatic gain control (AGC)," *Navigation, Journal of the Institute of Navigation*, vol. 59, no. 4, pp. 281–290, 2012.

[16] A. Cavaleri, B. Motella, M. Pini, and M. Fantino, "Detection of spoofed GPS signals at code and carrier tracking level," in *5th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing*, Dec. 2010.

[17] K. D. Wesson, B. L. Evans, and T. E. Humphreys, "A combined symmetric difference and power monitoring GNSS anti-spoofing technique," in *IEEE Global Conference on Signal and Information Processing*, 2013.

[18] A. J. Jafarnia, *GNSS Signal Authenticity Verification in the Presence of Structural Interference*. University of Calgary, 2013.

[19] D. S. D. Lorenzo, J. Gautier, J. Rife, P. Enge, and D. Akos, "Adaptive array processing for GPS interference rejection," in *Proceedings of the ION GNSS Meeting*. Long Beach, CA: Institute of Navigation, Sept. 2005.

[20] P. Y. Montgomery, T. E. Humphreys, and B. M. Ledvina, "A multi-antenna defense: Receiver-autonomous GPS spoofing detection," *Inside GNSS*, vol. 4, no. 2, pp. 40–46, April 2009.

[21] A. Broumandan, A. Jafarnia-Jahromi, V. Dehgahanian, J. Nielsen, and G. Lachapelle, "GNSS spoofing detection in handheld receivers based on signal spatial correlation," in *Proceedings of the IEEE/ION PLANS Meeting*. Myrtle Beach, SC: Institute of Navigation, April 2012.

[22] D. Borio, "PANOVA tests and their application to GNSS spoofing detection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 1, pp. 381–394, Jan. 2013.

[23] S. Daneshmand, A. Jafarnia, A. Broumandan, and G. Lachapelle, "GNSS spoofing mitigation in multipath environments using space-time processing," in *European navigation conference (ENC) 2013*, 2013, pp. 23–25.

[24] A. Konovaltsev, M. Cuntz, C. Haettich, and M. Meurer, "Performance analysis of joint multi-antenna spoofing detection and attitude estimation," in *Proceedings of the ION Internation Technical Meeting*, 2013.

[25] M. Psiaki, S. P. Powell, and B. W. O'Hanlon, "GNSS spoofing detection using high-frequency antenna motion and carrier-phase data," in *Proceedings of the ION GNSS+ Meeting*, 2013, pp. 2949–2991.

[26] S. Khanafseh, N. Roshan, S. Langel, F. Cheng-Chan, M. Joerger, and B. Pervan, "GPS spoofing detection using RAIM with INS coupling," in *Proceedings of the IEEE/ION PLANS Meeting*, May 2014.

[27] L. Scott, "Anti-spoofing and authenticated signal architectures for civil navigation systems," in *Proceedings of the ION GNSS Meeting*, 2003, pp. 1542–1552.

[28] K. D. Wesson, M. P. Rothlisberger, and T. E. Humphreys, "Practical cryptographic civil GPS signal authentication," *Navigation, Journal of the Institute of Navigation*, vol. 59, no. 3, pp. 177–193, 2012.

[29] T. E. Humphreys, "Detection strategy for cryptographic GNSS anti-spoofing," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 2, pp. 1073–1090, 2013.

[30] S. Lo, D. DeLorenzo, P. Enge, D. Akos, and P. Bradley, "Signal authentication," *Inside GNSS*, vol. 0, no. 0, pp. 30–39, Sept. 2009.

[31] M. L. Psiaki, B. W. O'Hanlon, J. A. Bhatti, and T. E. Humphreys, "Civilian GPS spoofing detection based on dual-receiver correlation of military signals," in *Proceedings of the ION GNSS Meeting*. Portland, Oregon: Institute of Navigation, 2011.

[32] M. Psiaki, B. O'Hanlon, J. Bhatti, D. Shepard, and T. Humphreys, "GPS spoofing detection via dual-receiver correlation of military signals," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 4, pp. 2250–2267, 2013.

[33] B. O'Hanlon, M. Psiaki, J. Bhatti, and T. Humphreys, "Real-time spoofing detection using correlation between two civil GPS receiver," in *Proceedings of the ION GNSS Meeting*. Nashville, Tennessee: Institute of Navigation, 2012.

[34] B. W. OHanlon, M. L. Psiaki, T. E. Humphreys, J. A. Bhatti, and D. P. Shepard, "Real-time GPS spoofing detection via correlation of encrypted signals," *Navigation, Journal of the Institute of Navigation*, vol. 60, no. 4, pp. 267–278, 2013.

[35] 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.

[36] K. D. Wesson, T. E. Humphreys, and B. L. Evans, "Receiver-autonomous GPS signal authentication based on joint detection of correlation profile distortion and anomalous received power," 2014, (in preparation).

[37] A. J. Kerns, K. D. Wesson, and T. E. Humphreys, "A blueprint for civil GPS navigation message authentication," in *Proceedings of the IEEE/ION PLANS Meeting*, May 2014.

[38] J. A. Bhatti and T. E. Humphreys, "Covert control of surface vessels via counterfeit civil GPS signals," 2014, (in preparation).

[39] A. J. Van Dierendonck, *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996, ch. 8: GPS Receivers, pp. 329–407.

[40] B. O'Hanlon, M. Psiaki, S. Powell, J. Bhatti, T. E. Humphreys, G. Crowley, and G. Bust, "CASES: A smart, compact GPS software receiver for space weather monitoring," in *Proceedings of the ION GNSS Meeting*. Portland, Oregon: Institute of Navigation, 2011, pp. 2745–2753.

[41] E. G. Lightsey, T. E. Humphreys, J. A. Bhatti, A. J. Joplin, B. W. OHanlon, and S. P. Powell, "Demonstration of a space capable miniature dual frequency GNSS receiver," *Navigation, Journal of the Institute of Navigation*, vol. 61, no. 1, pp. 53–64, 2014.

[42] T. Nighswander, B. Ledvina, J. Diamond, R. Brumley, and D. Brumley, "GPS software attacks," in *Proceedings of the 2012 ACM conference on Computer and communications security*. ACM, 2012, pp. 450–461.

[43] O. Montenbruck, A. Hauschild, and U. Hessels, "Characterization of GPS/GIOVE sensor stations in the CONGO network," *GPS Solutions*, vol. 14, no. 3, pp. 193–205, 2011.

[44] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, and P. M. Kintner, Jr., "GNSS receiver implementation on a DSP: Status, challenges, and prospects," in *Proceedings of the ION GNSS Meeting*. Fort Worth, TX: Institute of Navigation, 2006, pp. 2370–2382.

[45] T. E. Humphreys, J. Bhatti, T. Pany, B. Ledvina, and B. O'Hanlon, "Exploiting multicore technology in software-defined GNSS receivers," in *Proceedings of the ION GNSS Meeting*. Savannah, GA: Institute of Navigation, 2009, pp. 326–338.

[46] B. M. Ledvina, M. L. Psiaki, S. P. Powell, and P. M. Kintner, Jr., "Bit-wise parallel algorithms for efficient software correlation applied to a GPS software receiver," *IEEE Transactions on Wireless Communications*, vol. 3, no. 5, Sept. 2004.

[47] M. L. Psiaki, "Real-time generation of bit-wise parallel representations of over-sampled prn codes," *IEEE Transactions on Wireless Communications*, vol. 5, no. 3, pp. 487–491, March 2006.

[48] B. Ledvina, "Efficient real-time generation of bit-wise parallel representations of oversampled carrier replicas," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 47, no. 4, pp. 2921–2933, OCTOBER 2011.

[49] H. L. V. Trees, *Detection, Estimation, and Modulation Theory*. Wiley, 2001.

[50] H. V. Poor, *An Introduction to Signal Detection and Estimation, 2nd Edition*. Springer, 1994.

[51] T. E. Humphreys, M. L. Psiaki, and P. M. Kintner, Jr., "Modeling the effects of ionospheric scintillation on GPS carrier phase tracking," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 46, no. 4, pp. 1624–1637, Oct. 2010.

[52] D. Shepard and T. E. Humphreys, "Characterization of receiver response to a spoofing attack," in *Proceedings of the ION GNSS Meeting*. Portland, Oregon: Institute of Navigation, 2011.

[53] A. P. Cerruti, P. M. Kintner, D. E. Gary, A. J. Mannucci, R. F. Meyer, P. Doherty, and A. J. Coster, "Effect of intense December 2006 solar radio bursts on GPS receivers," *Space Weather*, vol. 6, no. 10, 2008.

[54] A. P. Cerruti, P. M. Kintner, D. E. Gary, L. J. Lanzerotti, E. R. de Paula, and H. B. Vo, "Observed solar radio burst effects on GPS/Wide Area Augmentation System carrier-to-noise ratio," *Space Weather*, vol. 4, no. S10006, Oct. 2006.

[55] G. Nita, D. Gary, L. Lanzerotti, and D. Thomson, "The peak flux distribution of solar radio bursts," *The Astrophysical Journal*, vol. 570, p. 423, 2002.

[56] T. E. Humphreys, *The GNSS Handbook*. Springer, 2014, ch. Interference, (in preparation).

[57] J. Do, D. M. Akos, and P. K. Enge, "L and S bands spectrum survey in the San Francisco Bay area," in *Proceedings of the IEEE/ION PLANS Meeting*. IEEE, 2004, pp. 566–572.

[58] T. E. Humphreys, "The GPS dot and its discontents: Privacy vs. GNSS integrity," *Inside GNSS*, vol. 7, no. 2, Mar./Apr. 2012.

[59] R. Mitch, R. Dougherty, M. Psiaki, S. Powell, B. O'Hanlon, J. Bhatti, and T. Humphreys, "Signal characteristics of civil GPS jammers," in *Proceedings of the ION GNSS Meeting*, 2011.