Characterization of Receiver Response to Spoofing Attacks

Daniel P. Shepard and Todd E. Humphreys
The University of Texas at Austin

BIOGRAPHIES

Daniel P. Shepard is pursing a Ph.D. in the Department of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin, where he also received his B.S. He currently works in the University of Texas at Austin Radionavigation Lab. His research interests are in GNSS security, estimation and filtering, and guidance, navigation, and control.

Todd E. Humphreys is an assistant 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. His research interests are in estimation and filtering, GNSS technology, GNSS-based study of the ionosphere and neutral atmosphere, and GNSS security and integrity.

ABSTRACT

Test procedures are developed for characterizing the response of civil GPS receivers to spoofing attacks. Two response characteristics are analyzed in detail for four representative GPS receivers: (1) the spoofer power advantage over the authentic signals required for successful receiver capture, and (2) the aggressiveness with which a spoofer can manipulate the victim receiver's time and position solution. Two of the tested receivers are commonly used in critical infrastructure applications, one in smart power grid regulation and one in telecommunications networks. The implications of the test results for these critical infrastructure applications are discussed.

I. INTRODUCTION

In 2001, the U.S. Department of Transportation (USDOT) evaluated the transportation infrastructure's vulnerability to GPS and raised concern over the threat of GPS spoofers [1]. Spoofers generate counterfeit GPS signals that commandeer a victim receiver's tracking loops and induce spoofer-controlled time or position offsets. The USDOT report noted the absence of any off-the-shelf defense against civilian spoofing and recommended a study to characterize spoofing effects and observables. In 2008, researchers demonstrated that an inexpensive portable software-defined GPS spoofer could be built from off-the-shelf components, again highlighting the threat of spoofing [2].

Cell phone networks are one segment of critical infrastructure that is vulnerable to civil spoofing attacks. CDMA cell phone towers rely on GPS timing for tower-to-tower synchronization. Synchronization prevents the towers from interfering with one another and enables call hand-off from one tower to the next. If a particular tower deviates more than $10\mu s$ from GPS time, hand-off to and from that tower is disrupted and overall network throughput is reduced [3].

The power grid also possesses a unique vulnerability to spoofing attacks. More efficient distribution of power across the grid will require real-time measurements of the voltage and current phasors [4]. Synchrophasor Measurement Units (SMUs) have been proposed as a smart grid technology for precisely this purpose. SMUs rely on GPS to time stamp their measurements, which are sent back to a central monitoring station for processing. Manipulation of a SMU's time stamp results in spurious variations in the measured phase angles. These variations could distort power flow or stability estimates in such a way that grid operators would take incorrect or unnecessary control actions including powering up or shutting down generators. This could cause blackouts or damage to power grid equipment. Reference [5] also discusses the effects that alteration of SMU time stamps might have on fault location. This could hamper the ability of utility companies to respond quickly and appropriately to faults.

While much promising research is currently being conducted on methods for detecting and mitigating spoofing, these methods are still years away from wide-scale implementation. Meanwhile, it is important to understand what risks a spoofing attack poses for existing GPS receivers and for critical infrastructure reliant on civil GPS for positioning, timing, or both. The specific questions that this paper seeks to answer through experimentation are:

1. Would a jamming-power-to-noise-power (J/N)-type detector, which is commonly used in military GPS receivers to detect jamming attacks, trigger on a spoofing attack?

2. How aggressively can a civil GPS receiver's navigation and timing solution be manipulated by a spoofing attack?

The four receivers tested are meant to be a representative cross-section of civil user equipment. The receivers are (1) a science-grade receiver, (2) a time reference receiver with a highly stable internal clock used in telecommunications

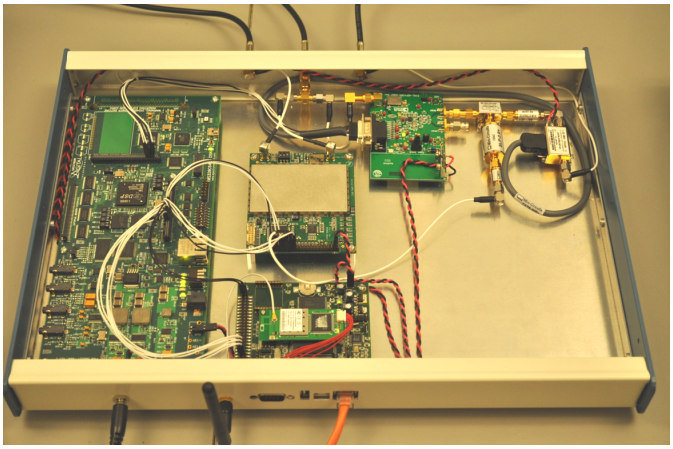

Fig. 1. The Civil GPS Spoofer.

networks, (3) a time reference receiver with a less stable internal clock used on the power grid, and (4) a consumer-grade handheld receiver.

II. THE SPOOFER

The Civil GPS Spoofer used for these tests, shown in Fig. 1, is an advanced version of the spoofer reported in [2]. It is the only spoofer reported in open literature to date that is capable of precisely aligning the spreading code and data bits and matching the frequency of its counterfeit signals with the authentic GPS signals. Such alignment capability allows the spoofer to carry out a sophisticated spoofing attack in which no obvious clues remain to suggest that an attack is underway. This spoofer is implemented on a portable software-defined radio platform with a digital signal processor (DSP) at its core. This platform comprises:

- A Radio Frequency (RF) front-end that down-mixes and digitizes GPS L1 and L2 frequencies
- A DSP board that performs acquisition and tracking of GPS L1 C/A and L2C signals, calculates a navigation solution, predicts the L1 C/A databits, and produces a consistent set of up to 10 spoofed GPS L1 C/A signals with a fictitious implied navigation and timing solution.
- An RF back-end with a digital attenuator that converts the digital samples of the spoofed signals from the DSP to analog output at the GPS L1 frequency with a user-controlled broadcast power.
- A Single Board Computer (SBC) that handles communication between the spoofer and a remote computer over the Internet.

The spoofer works by first acquiring and tracking GPS L1 C/A and L2C signals to obtain a navigation solution. Once a navigation solution has been obtained, the spoofer enters its "feedback" mode. In this mode, the spoofer produces a counterfeit, data-free feedback GPS signal that is summed with its own antenna input. The feedback signal is tracked

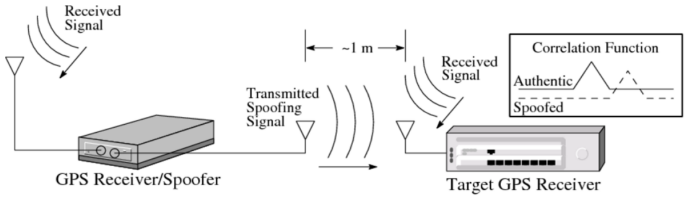

Fig. 2. A proximity spoofing attack [2].

by the spoofer and used to calibrate the delay between production of the digitized spoofed signal and output of the analog spoofed signal. This is necessary because the delay is non-deterministic on start-up of the receiver, although it stays constant thereafter.

After feedback calibration is complete and enough time has elapsed to build up a databit library, the spoofer is ready to begin an attack. It produces signals that are initially nearly perfectly aligned with the authentic signals at a low power to remain below the noise floor. The spoofer then raises the power of the spoofed signals slightly above that of the authentic signals. At this point, the spoofer has taken control of the victim receiver's tracking loops and slowly leads the spoofed signals away from the authentic signals, carrying the receiver's tracking loops with it. Once the spoofed signals have moved more than 600 m in position or 2 μs in time away from the authentic signals, the receiver has been completely captured.

For the experiments reported in this paper, the spoofed signals were not broadcast over the air, but were routed via coaxial cable to the antenna input of the target receiver, where they were summed with the authentic signals. This configuration is representative of a proximity spoofing attack, where the spoofer is within a meter or so of the target receiver, as shown in Fig. 2. In this case, the spoofer does not need to account for the distance between its antenna and the target receiver's antenna. This type of attack was originally envisioned in [6].

III. APPROACH

The following two subsections reduce the two questions posed in the Introduction to specific component questions that can be addressed by experimentation.

A. Jamming Detector

A J/N-type jamming detector works by comparing the total received in-band power with the in-band noise power, or the in-band power as measured by a GPS receiver and antenna combination in the absence of external signals (e.g. within an anechoic chamber). Account is taken of natural variations in in-band power due to satellite geometry and solar activity. To avoid frequent false alarms, a J/N-type jamming detector triggers only if the measured J/N ratio exceeds a threshold above which natural variations only rarely push the J/N measurement. In [7] it is estimated

that such a trigger would be insensitive to civil GPS spoofing attacks up to a spoofing power ratio of 3, where the spoofing power ratio is defined as

$$\eta = \frac{P_{spoof}}{P_{auth}} \tag{1}$$

with P_{spoof} being the power of the spoofing signal ensemble and P_{auth} being the power of the authentic signal ensemble. If all authentic signals are spoofed, then η is also the ratio of each individual spoofing signal to its authentic counterpart. The question that needs to be answered through testing is this: What power ratio is required for reliable spoofing?

B. Aggressive Receiver Manipulation

It is important to understand the types of dynamics that a spoofer could induce in a target receiver. For example, SMU measurements of power grid stability could be manipulated by spoofer-induced oscillations in the SMU's time reference if the SMU were susceptible to extremely aggressive GPS spoofing.

Three specific questions regarding spoofer induced dynamics are posed:

1. How quickly could a timing or position bias be introduced?

2. What kinds of oscillations could a spoofer cause in a receiver's position and timing?

3. How different are receiver responses to spoofing?

The approach taken to answer these questions was to determine the maximum velocities that can be induced in a target receiver's position or time solution over a range of accelerations when starting from a velocity of zero. The curve in the velocity-acceleration plane created by connecting these points defines the upper bound of a region in which the spoofer can safely manipulate the target receiver without raising any alarms or causing the target receiver to lose satellite lock. Figure 3 shows four conceivable shapes for this curve (a vertical line, a horizontal line, a line with a negative slope, and an exponential curve) with the green area representing the safe region for a spoofer to operate and the red area representing the region where a spoofer will likely be caught. Once these curves have been obtained, they can be used to determine the kinds of dynamics a spoofer could induce on that receiver and can be compared to curves of other receivers to find which receivers are most resistant to a spoofing attack.

IV. PROCEDURE

Four receivers were tested in laboratory experiments:

Science receiver: The CASES receiver developed by the UT Radionavigation Lab in collaboration with Cornell University and ASTRA. This receiver was designed for ionospheric monitoring [8].

Telecommunications network time reference receiver: The HP 58503B, which has been commonly used in cell phone base stations. The 58503B has a highly stable oven controlled crystal oscillator (OCXO) steered by the GPS time solution [9].

Power grid time reference receiver: The SEL-2401, which provides the time signal for most power grid Synchrophasor Measurement Units (SMUs). SMUs are a proposed smart grid technology to make real time measurements of the voltage phasor (a 60 Hz phasor) for stability analysis, power flow estimation, and fault location on the power grid for monitoring and control purposes. It has a low stability oscillator (most likely a temperature controlled crystal oscillator (TCXO) or simple crystal oscillator (XO)) slaved to the GPS time solution [10].

Name brand receiver: The Trimble Juno SB, a high quality handheld receiver [11].

Pictures of all of these receivers are shown in Fig. 4. These receivers are meant to be a representative cross-section of civil GPS receivers.

Tests were performed in a controlled signal environment with a set of six GPS L1 C/A signals generated by a National Instruments Radio Frequency Signal Generator (RFSG) at a constant power level. Limiting the signals to a set of six GPS L1 C/A signals at constant power simplified the tests without significantly affecting the generality of the results. The signals were tracked by the spoofer, which produced a set of six corresponding spoofed signals. The spoofed signals were then summed with the RFSG-generated signals. This combination of spoofed and RFSG-generated signals was fed into both the target receiver and a National Instruments Radio Frequency Signal Analyzer (RFSA). The RFSA was used for visualization of the spoofing attack and measurement of the signal power. Figure 5 shows the described test setup.

A. Power Ratio Test

The power ratio test was performed by first setting the digital attenuator on the spoofer's RF back-end so that the spoofed signals were slightly stronger than the authentic signals with the signals still aligned. The spoofer then attempted to capture the target receiver by advancing the spoofed signals time solution forward at a rate of 3.33 ns/s, or about 1 m/s in equivalent velocity units.

Once every spoofed signal's correlation peak had completely separated from its counterpart authentic signal's correlation peak (a separation greater than 2 μs), the authentic signals were removed. The target receiver's output was then observed to see if any signals were lost. If no signals were lost, then the spoofing attack was deemed successful in capturing the receiver. The authentic signals

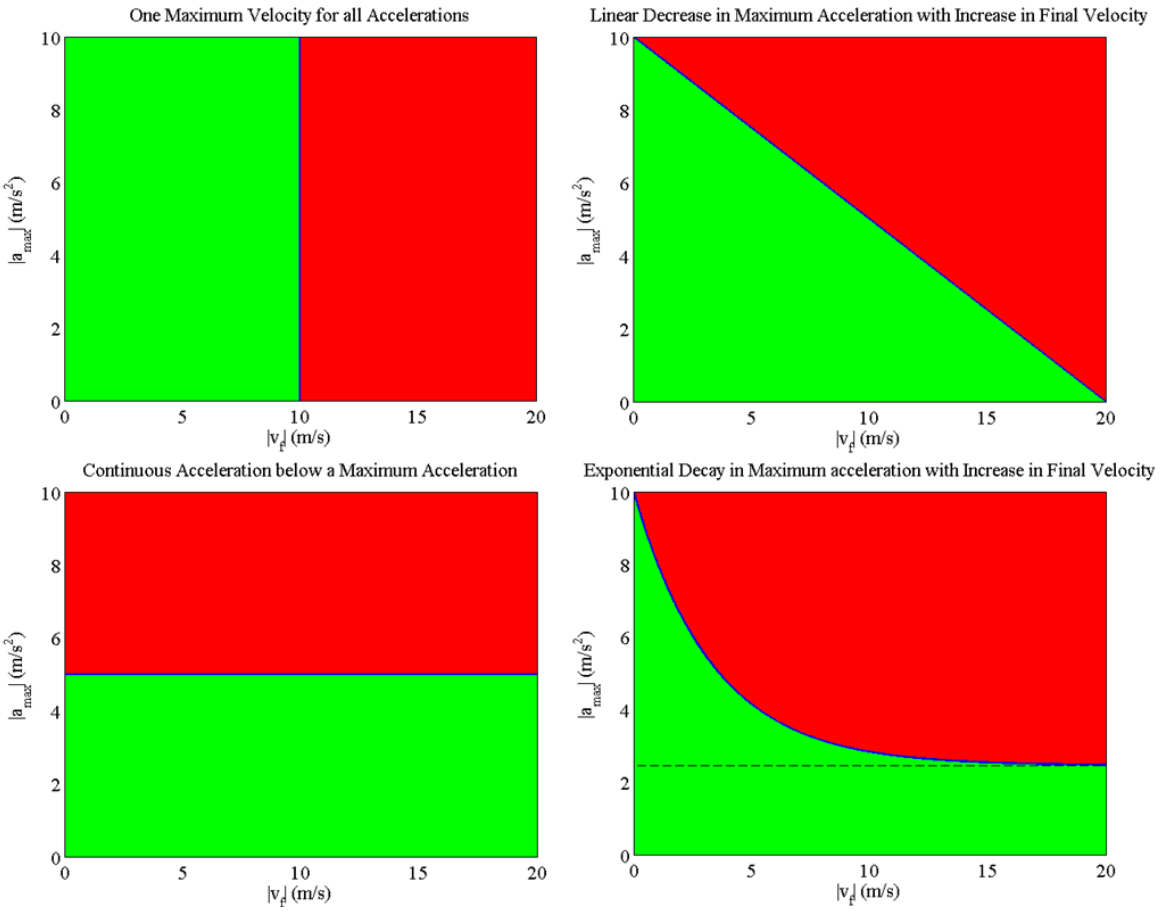

Fig. 3. Plots of four potential shapes of maximum acceleration-velocity curves for receiver position or time dynamics that can be induced by a spoofer. The green area represents the region where a spoofer can safely operate, and the red area represents the region where a spoofer will likely be caught. Time rate and acceleration are expressed in equivalent m/s and m/s^2.

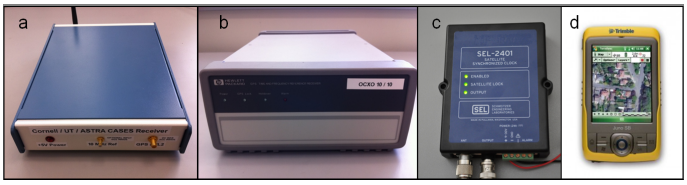

Fig. 4. The tested receivers: (a) CASES receiver, (b) HP 58503B, (c) SEL-2401, and (d) Trimble Juno SB.

were then reinserted and the relative power of each authentic and counterfeit signal was measured using the RFSA. After recording the results, this process was repeated a number of times. The spoofer's attenuator setting was modified as needed to find a power ratio limit above which a spoofer could consistently capture the target receiver. Figure 6 summarizes this procedure in a graphical format.

B. Spoofed Acceleration and Velocity Test

The spoofed acceleration and velocity test was performed with only the spoofing signals as input to the target receiver to simulate the behavior of the receiver after it has been captured. This approach eliminated interference with the authentic signals during these tests. First, an acceleration was set for the spoofer. Next, a final velocity was

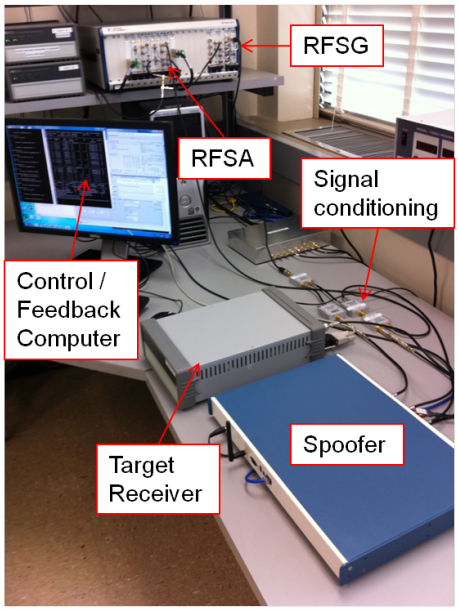

Fig. 5. The experimental setup.

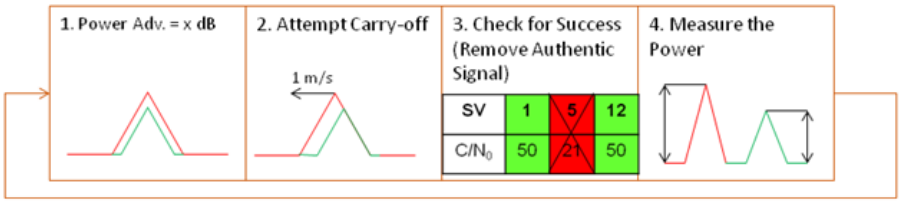

Fig. 6. Graphical representation of the procedure for the power ratio test.

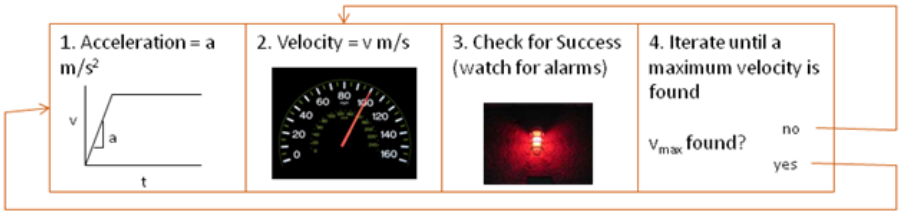

Fig. 7. Graphical representation of the procedure for the spoofed velocity and acceleration test.

chosen for the spoofer to reach. The output from the target receiver was monitored for alarms or loss of lock on any satellites until the receiver stabilized at the final velocity. If any evidence of the spoofing attack was evident in the receiver alarms or tracking status, then the spoofing attack was deemed unsuccessful. This process was repeated while modifying the final spoofed velocity until a maximum spoofed velocity was determined for that acceleration. Once this maximum velocity was found, the spoofed acceleration was modified and the process repeated until 5 or 6 data points were collected. Figure 7 summarizes this procedure in a graphical format.

V. RESULTS

A. Spoofing Power Ratio Test

The power ratio test results are summarized by the histogram shown in Fig. 8. This histogram shows the number of successful and failed spoofing attacks for the range of power ratios tested. This test was performed on both the science receiver and the telecommunications time reference receiver. Since both receivers displayed similar results, it was deemed unneccesary to conduct this test with the other two receivers.

As can be seen from Fig. 8, it is possible to reliably spoof a target receiver at a power ratio of 1.1. This value of η would keep the received in-band power well below the natural variations due to constellation changes and solar activity. This means that a J/N-type jamming detector would not necessarily detect a spoofing attack.

However, a J/N-type jamming detector is still an important component in spoofing detection schemes. One such scheme, the Vestigial Signal Defense (VSD) [2] [12], involves detecting the vestige of the authentic signal and distinguishing it from a multipath signal, which can only be done if the authentic signal has not been drowned out

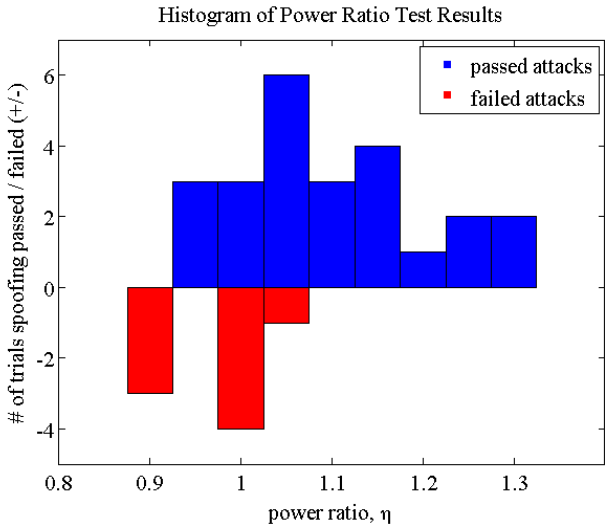

Fig. 8. Histogram of the spoofing power ratio test results performed on the science receiver and the telecommunications time reference receiver.

or nulled by the spoofer. Nulling the authentic signal is inherently difficult since it requires precise anti-alignment of the phase of a spoofed replica signal with the authentic signal. This makes drowning out the authentic signal the more likely attack scenario. A J/N type jamming detector would constrain the spoofer to operate with a power ratio below about 3, effectively eliminating the possibility of the spoofer entirely suppressing the authentic signal with excessive spoofing power.

B. Spoofed Velocity and Acceleration Test

The velocity and acceleration data points collected for each of the four tested receivers were fit to a single curve for each receiver of the form:

$$|a_{max}| = f(v_f) = \beta_1 e^{-\beta_2 |v_f|} + \beta_3 \qquad (2)$$

TABLE I

RAW DATA FOR THE SCIENCE RECEIVER

acceleration m/s^2	velocity m/s
8.8	2.2
7	4.6
6	6.4
5.5	8
5	1300

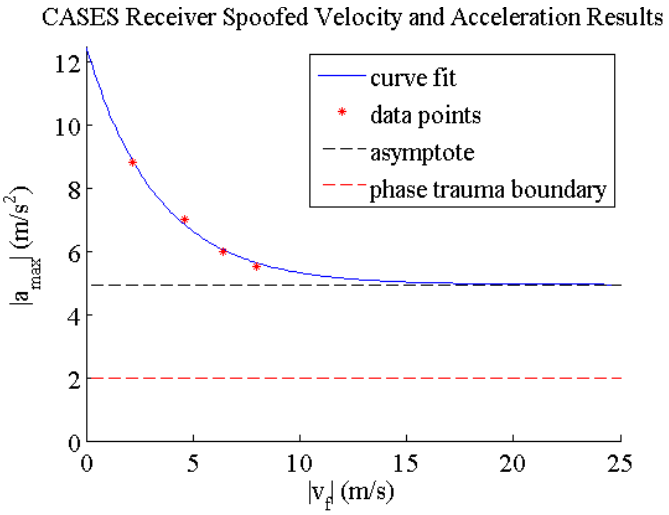

Fig. 9. Spoofed velocity and acceleration curve fit for the science receiver.

where v_f is the final velocity of the spoofer, a_{max} is the maximum acceleration the spoofer can use to reach the final velocity without triggering alarms or causing dropped signals, and β_1, β_2, and β_3 are fit parameters. An exponential model was chosen for this data based on intuition and on the testing results themselves. This model is meant to represent the capability of the receiver's discrete time tracking loops to remain locked to the GPS signal under the spoofer imposed dynamics. It defines the upper bound of a region of the velocity-acceleration plane in which the spoofer can safely operate. Knowledge of this curve for a particular receiver allows one to assess the security implications of a spoofing attack on a system dependent on the receiver.

B.1 Science Receiver (CASES)

The raw velocity and acceleration data points for the science receiver are given in Table I. These data were fit to Eq.(2) and plotted along with the curve fit in Fig. 9. The resulting values for the fit parameters are listed in Table II.

The first interesting feature to note in Fig. 9 is that there is a horizontal asymptote at an acceleration of about 5 m/s^2. This suggests that the science receiver can be accelerated continuously at accelerations below 5 m/s^2. The only limit to the velocity that can be induced in the science receiver is due to the Doppler frequency range of the receiver. Outside this range, the receiver is unable to produce local carrier replicas at the appropriate frequency resulting in a loss of satellite lock. This Doppler range is set to $\pm$ 10,000 Hz for stationary applications. The maximum attainable velocity offset is thus dependent on the exact satellite geometry, but is generally around 1,300 m/s. Beyond this speed, the receiver fails to track some satellites due to the large Doppler.

TABLE II

FIT PARAMETERS FOR THE SCIENCE RECEIVER

β_1	β_2	β_3
7.55	0.3	4.94

TABLE III

RAW DATA FOR THE TELECOMMUNICATIONS NETWORK TIME REFERENCE RECEIVER

acceleration m/s^2	velocity m/s
5.8	1.45
5	1.7
4.5	1.8
4	1.9
2	2
1	2

Another interesting feature of this receiver's response to induced dynamics is that above accelerations of about 2 m/s^2 the receiver indicates a constant state of phase trauma. The receiver's phase trauma flag indicates that its Phase Lock Loop (PLL) may be experiencing cycle slips. This is a special feature of the science receiver used to measure the effects of ionospheric scintillation [13]. Indications of phase trauma in the absence of scintillation could be considered an alarm. This effectively limits a spoofer that is trying to influence the navigation solution of the science receiver to accelerations below 2 m/s^2.

B.2 Telecommunications Network Time Reference Receiver (HP)

The raw velocity and acceleration data points for the telecommunications network time reference receiver are given in Table III. These data were fit to Eq.(2) and plotted along with the curve fit in Fig. 10. The resulting values for the fit parameters are listed in Table IV.

As can be seen from a comparison of Fig. 9 and Fig. 10, the HP receiver is much more resistant to dynamics than the science receiver. The maximum velocity that could be successfully induced in the telecommunications network time reference receiver was 2 m/s. This resistance to spoofer induced dynamics is due to the trust that the HP receiver places in its highly stable oscillator. The receiver was de-

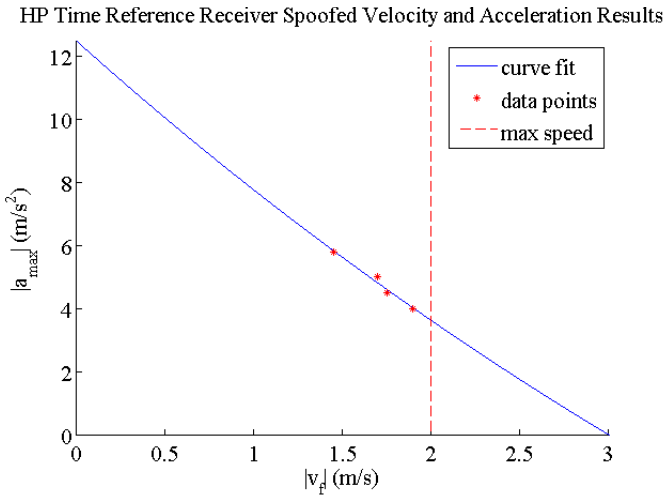

Fig. 10. Spoofed velocity and acceleration curve fit for the telecommunications network time reference receiver.

TABLE IV
FIT PARAMETERS FOR THE TELECOMMUNICATIONS NETWORK TIME REFERENCE RECEIVER

β_1	β_2	β_3
38.02	1.33e-1	-25.53

signed such that its oscillator's time output is only loosely coupled to the GPS time solution. The oscillator is slowly steered into alignment with the GPS time solution. The receiver enters a "holdover" mode if the difference between the GPS time solution and the receiver's oscillator time is greater than 1 μs. This feature acts as an alarm to indicate that GPS should no longer be trusted. It is this feature that causes the deviation from the curve fit in Fig. 10. There is simply no spoofer acceleration that will allow the HP receiver to stabilize at a speed greater than 2 m/s before entering holdover mode.

B.3 Power Grid Time Reference Receiver (SEL-2401)

The raw velocity and acceleration data points for the power grid time reference receiver are given in Table V. These data were fit to Eq.(2) and plotted along with the curve fit in Fig. 11. The resulting values for the fit parameters are listed in Table VI.

As can be seen in Fig. 11, the power grid time reference receiver can be manipulated fairly easily by a spoofer. A secondary x-axis was added to the figure to represent the corresponding phase angle rate that could be induced in a 60-Hz phasor being measured using the receiver's time output. These results showed that the power grid time reference receiver can reach a maximum speed of 400 m/s, which corresponds to a 1.73 ^{o}/min phase angle rate for a voltage phasor on the power grid. At speeds above 400 m/s, the receiver sporadically looses and regains lock on the satellites, which could be considered an alarm.

TABLE V
RAW DATA FOR THE POWER GRID TIME REFERENCE RECEIVER

acceleration m/s^2	velocity m/s
10	2.5
8	75
7	120
5	190
3	360
2	400
1	400

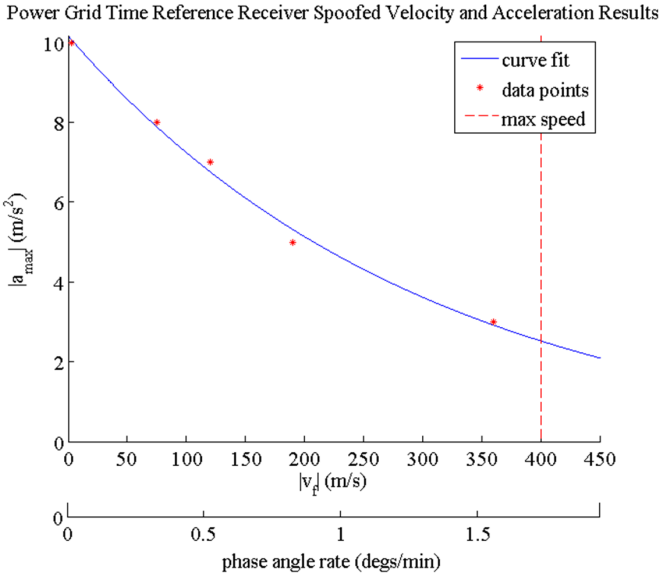

Fig. 11. Spoofed velocity and acceleration curve fit for the power grid time reference receiver with secondary x-axis corresponding to the induced phase angle rate for a 60 Hz phasor (such as the voltage phasor for the power grid).

B.4 Name Brand Receiver (Trimble)

The raw velocity and acceleration data points for the name brand receiver are given in Table VII. These data were fit to Eq.(2) and plotted along with the curve fit in Fig. 12. The resulting values for the fit parameters are listed in Table VIII.

From comparisons of Fig. 12 with Figs. 9, 10, and 11, the Trimble is by far the most easily manipulated receiver of those tested. There is a horizontal asymptote at around 25 m/s^2, which suggests that the Trimble can be accelerated continuously at accelerations below 25 m/s^2. As with the science receiver, the only limit to the velocity that can be induced is due to the Doppler frequency range of the receiver. This Doppler range appears to be about $\pm$ 10,000 Hz based on the result that the receiver fails to track some satellites at velocities greater than 1,300 m/s, which was also the case for the science receiver. These results suggest that the Trimble's robustness – its ability to provide navigation and timing solutions despite extreme signal dynamics – is actually a liability in regard to spoofing. In

TABLE VI

FIT PARAMETERS FOR THE POWER GRID TIME REFERENCE RECEIVER

β_1	β_2	β_3
10.48	3.3e-3	-0.31

TABLE VII

RAW DATA FOR THE NAME BRAND RECEIVER

acceleration m/s^2	velocity m/s
49.2	12.3
35	15.7
30	19.5
27	23
25	190
24.5	1300

other words, the Trimble receiver's ability to track high accelerations and velocities allows a spoofer to aggressively manipulate its outputs.

VI. Implications

A. Implications for Cellular CDMA Communications Networks

Although the HP receiver provides a high resistance to imposed dynamics, it can still be led off far enough in time to cause some harmful effects on CDMA cell phone networks that typically use such receivers. A 10 μs time offset can be imparted in around 35 minutes, including time for receiver capture. At this time offset the dependent cell phone tower becomes an "island tower," unable to transfer calls to and from adjacent towers [3]. At larger time offsets, approaching 60 μs, the tower's signal could begin to interfere with signals from neighboring towers, causing a further disruption in the network. This is because CDMA cell phone towers all use the same spreading code and distinguish themselves only by start time of the code.

Spoofer-induced time offsets might also be a problem for other uses of the cell phone network, such as E911-style localization [14]. E911 or Enhanced 911 is a system that uses Time Difference Of Arrival (TDOA) techniques to locate cell phone users who dial 911 in emergency situations.

B. Implications for Power Grid Monitoring and Control

B.1 State Estimation

One of the most important uses of SMUs for the smart grid is real-time state estimation. It has been suggested that SMUs are necessary to provide accurate, real-time

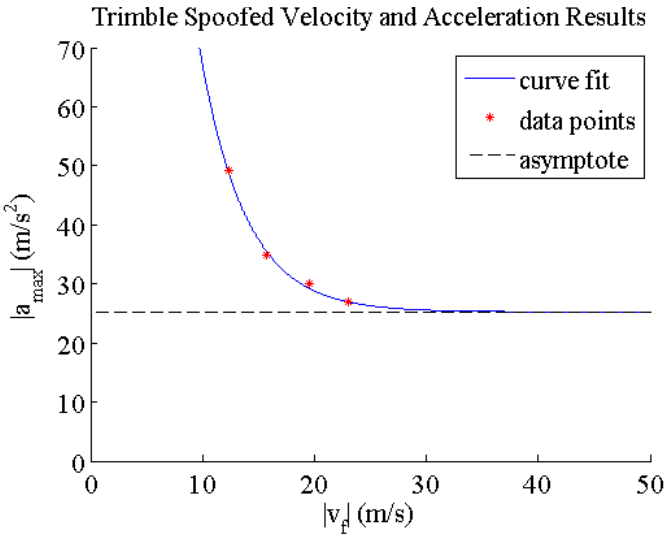

Fig. 12. Spoofed velocity and acceleration curve fit for the name brand receiver.

TABLE VIII

FIT PARAMETERS FOR THE NAME BRAND RECEIVER

β_1	β_2	β_3
444.16	0.24	24.92

estimates of the state of the power grid so that power margins can be reduced to make the grid more efficient [4]. Spoofing poses a risk to state estimation because a change in the voltage phase angle difference between two locations in the grid directly relates to a change in the estimated power flow between those locations. Alteration of the voltage phase angle at a particular location by even 10° could cause an operator or automatic control logic to take incorrect or unnecessary control action. A spurious 10° phase deviation can be imposed by the spoofer in a matter of minutes.

Such a large variation in the phase angle difference may not be seen as suspicious because wind power generation leaves behind similar signatures. This can be seen clearly thanks to a proof of concept network set up on the Texas power grid [15]. Figure 13 shows the variation of wind power generation in west Texas on March 10, 2009. There is a large spike and a subsequent drop in wind power generation during the 11:00pm to midnight hour. The SMU data, shown in Fig. 14, reveals a corresponding rise in the phase angle difference between Austin and west Texas. A grid operator looking at SMU measurements would not be able to tell the difference between a spike in wind generation and a spoofing attack without the benefit of direct measurements of the wind power production. This suggests that current power flow meters, if retained and monitored, could provide a cross-check for SMU power flow estimates.

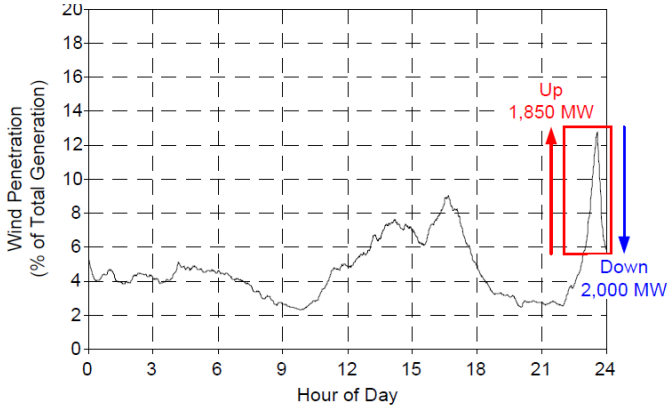

Fig. 13. Texas wind generation on March 10, 2009 [15]. Used with permission from Mack Grady.

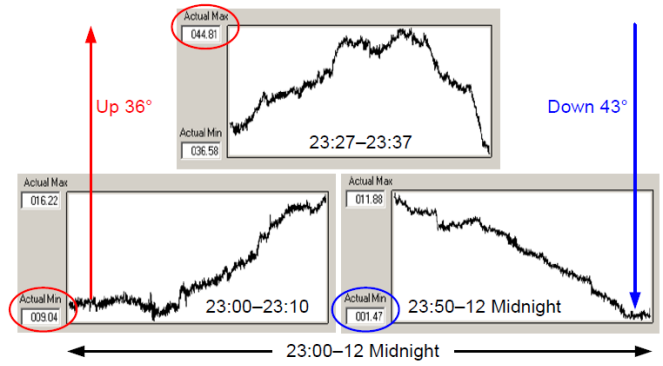

Fig. 14. SMU measured voltage phase angle difference between Austin and west Texas during wind generation spike on March 10, 2009 [15]. Used with permission from Mack Grady.

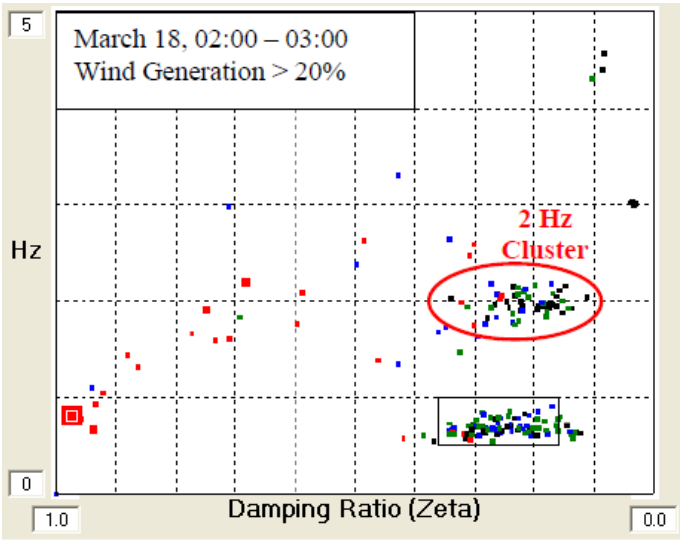

Fig. 15. Example of SMU calculated damping ratios and frequencies over an hour long period. The color of the dots indicate the magnitude of the oscillation with the largest 25% marked red, the second 25% marked blue, the third 25% marked green, and the lowest 25% marked black. In this case, the red are several degrees in magnitude and all others are less than a degree [15]. Used with permission from Mack Grady.

B.2 Stability Determination

Another important application of SMUs is the determination of stability of the power grid. Unstable, low frequency oscillations in the voltage phase angle can damage power generators if no corrective action is taken. These low-frequency oscillations occur due to changes in the load or power generation on the power grid with the amplitude of the oscillations scaling with the magnitude of the load or power generation. Most often these oscillations are damped by the power system stabilizers on the generators, but larger disturbances can be difficult for these stabilizers to handle on their own and the damping could become negative. The oscillations of concern are of magnitudes greater than several tenths of a degree and frequencies between 0.1 Hz and 0.8 Hz [15].

Oscillations on the power grid are modeled as the superposition of multiple second-order systems. The frequencies and damping coefficients of these oscillations are estimated using a modified version of the Prony Method [16], a close cousin of the Fourier Transform that works on discrete data points taken from a moving window. The data points are modeled as a linear combination of damped sinusoids and fit to a function of the form

$$f[mt_s + t_0] = \sum_{i=1}^{n} \frac{1}{2} A_i e^{(\sigma_i \pm j(\lambda_i(mt_s+t_0)+\phi_i))} \qquad (3)$$

where m is the sample number, t_s is the sampling interval, t_0 is the start of the window, $f[mt_s + t_0]$ is the sample, n is the number of damped sinusoids used for the fit, and A_i, σ_i, λ_i, and ϕ_i are the amplitude, damping coefficient, frequency, and phase angle of the ith damped sinusoid.

Figure 15 shows a plot of the observed damping ratios on the Texas power grid calculated using this method over an hour-long period. There is a large cluster of points, indicating a persistent mode of the system, at 0.7 Hz and a 2-Hz cluster that appears due to high wind power generation [15]. There are also a number of high amplitude oscillations, indicated by red dots, that appear occasionally during this time frame. These points have a large damping ratio, likely due to the power system stabilizers, and quickly die off.

Based on these tests, it seems impossible for a spoofer to cause oscillations in the PMU measurements of sufficient magnitude at an appropriate frequency to affect power grid stability estimates. This is due to the low acceleration capability of the spoofer. An oscillation at 0.1 Hz with an amplitude of 0.1^o would require a maximum acceleration of about 550 m/s^2 and a maximum velocity of about 900 m/s. These values are far beyond the admissible dynamics indicated in Fig. 11.

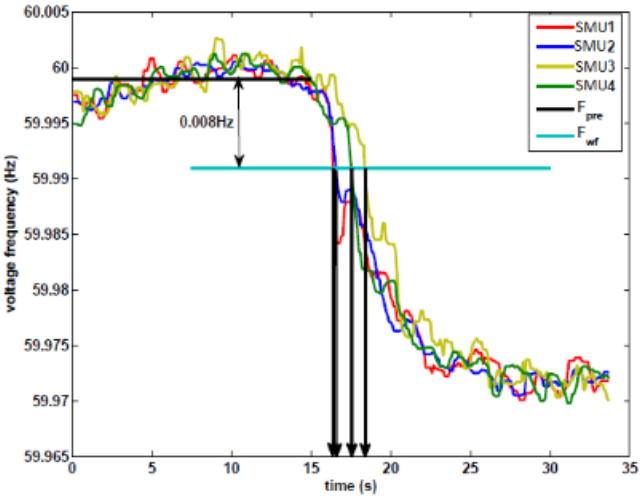

Fig. 16. Time of arrival determination using the leading edge of a frequency fault as seen by multiple PMUs at different locations [5]. Used with permission from Qi Zeng.

B.3 Fault Localization

Another possible use of PMUs is fault localization. One manifestation of a fault is a sudden drop or rise in frequency. These effects propagate through the power grid at the speed of light and will be observed by PMUs at different locations at different times depending on the distance from the source of the fault. Fig. 16 shows time of arrival determination using the leading edge of the fault as it arrives at different PMUs [5]. Using this information, Time Difference Of Arrival (TDOA) techniques can be used to locate the origin of the fault and remedial actions can be taken.

GPS spoofers could potentially alter the time stamp of the detected fault on one or more of the PMUs. This would corrupt the estimate of the fault origin and hamper repair efforts. Reference [5] investigates the results that one might obtain if such an attack were carried out. In one particular scenario from Ref. [5], the spoofer alters the timing such that a fault that occurred in Indiana appears to have occurred in Alabama.

VII. CONCLUSIONS

Results of tests designed to characterize the response of a civil GPS receiver to a spoofing attack indicate that a J/N-type jamming detector is insufficient to catch a spoofer. The ratio of spoofed signal power to authentic signal power required to consistently capture a target receiver is only about 1.1. This increase in J/N would typically be ignored by a jamming detector because it is within the natural variation in J/N caused by GPS satellite constellation changes and solar activity [7]. However, a J/N type jamming detector is an essential component in many potential spoofing defenses, including the Vestigial Signal Defense (VSD), since it limits the amount of power a spoofer can surreptitiously transmit, which prevents the spoofer from completely suppressing the authentic signal by making it appear as noise.

Investigations into the dynamics that a spoofer can induce in a target receiver yielded results that varied drastically between four tested receivers. An empirical curve fit for the maximum acceleration that can be used by a spoofer to reach a certain final velocity in position or timing without raising alarms or causing loss of satellite lock in the target receiver was produced for each receiver based on an exponential model. These curve fits define the upper bound of a region of the velocity-acceleration plane in which the spoofer can operate without triggering alarms or causing loss of satellite lock in the target receiver. These empirical formulas can be used to assess the vulnerability of critical infrastructures utilizing these receivers.

The science receiver results showed that there is no limit to the velocity a spoofer can induce in the receiver until the Doppler frequency range of the receiver is exceeded at around 1,300 m/s. However, the acceleration was severely limited due to the science receiver's constant indication of phase trauma during accelerations above 2 m/s^2, which can be treated as an alarm in the absence of scintillation. This type of visibility into the receiver's tracking loops provides an advantage towards limiting a spoofer's capability of dynamically manipulating a receiver.

The telecommunications network time reference receiver was by far the most difficult receiver to manipulate, with a maximum attainable velocity for any acceleration being only 2 m/s. This receiver's inherent resistance to spoofer imposed dynamics is due to the receiver placing trust in its oscillator and only slowly steering it towards the GPS time solution. However, the receiver can still be slowly steered away from GPS time: a 10 μs departure can be forced in around 35 minutes, including time for capturing the receiver. This time offset is enough to prevent call hand-off to or from a CDMA cell phone tower. A spoofer – or spoofer network – could also cause multiple neighboring towers to interfere with one another, since CDMA cell phone towers all use the same spreading code and distinguish themselves only by the phasing (i.e. time offset) of their spreading codes. Furthermore, it appears possible for a spoofer to impair CDMA-based E911 user-location.

The power grid time reference receiver was fairly easily manipulated, with a maximum attainable velocity of 400 m/s, but it could only track single-digit accelerations. This receiver is typically used as the time reference for Synchrophasor Measurement Units (SMUs), which measure voltage phasors on the power grid. SMUs are a proposed smart grid technology that will provide real-time stability analysis, power flow state estimation, and fault localization. A spoofer could easily cause large variations in the power flow estimates from SMU data by altering the receiver's time stamp, which in turn changes the voltage phase angle suggesting a change in the power flow. The maximum attainable phase angle rate from these tests was

1.73 ^{o}/min. These changes in power flow measurements could cause a grid operator or automatic control logic to take corrective actions based on falsified data, potentially resulting in damage to the power grid. Current power flow meters could provide a valuable cross-check against the SMU derived power flow estimates. In order to affect grid stability measures, a spoofer would be required to falsify unstable, low-frequency oscillations in the phase measurements. Based on the test results, it appears that a spoofer is incapable of producing such oscillations at the appropriate frequencies with sufficient magnitudes. A spoofer might also affect fault location estimates obtained through time difference of arrival (TDOA) techniques using SMU measurements. This could potentially cause large errors in these location estimates which would hamper repair efforts.

The name-brand receiver was by far the easiest receiver to manipulate, with continuous acceleration possible up to 25 m/s^2, and velocity is only limited by the Doppler frequency range of the receiver. This suggests that the navigation and timing solution of portable GPS receivers meant to operate under a wide variety of platform dynamics could be aggressively manipulated by a spoofer.

ACKNOWLEDGMENTS

The authors thank Mack Grady for setting up the Texas Synchrophasor Network and providing data from the network and the University of Texas at Austin Radionavigation Laboratory for their aid in developing the spoofer.

References

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

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

[3] "Recommended Minimum Performance Standards for cdma2000 Spread Spectrum Base Stations, C.S0010-B," Tech. rep., 3rd Generation Partnership Project 2 "3GPP2", Feb. 2004.

[4] Giri, J., Sun, D., and Avila-Rosales, R., "Wanted: A more intelligent grid," *IEEE Power & Energy*, April 2009, pp. pp. 34–40.

[5] Zhang, Z., Gong, S., Li, H., Pei, C., Zeng, Q., and Jin, M., "Time stamp attack on wide area monitoring system in smart grid," *Computing Research Repository*, Feb. 2011.

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

[7] Humphreys, T. E., Shepard, D., and Bhatti, J., "A Testbed for Developing and Evaluating GNSS Signal Authentication Techniques," *IEEE Transactions on Aerospace and Electronic Systems*, 2011, in preparation; available at http://radionavlab.ae.utexas.edu/testbed.

[8] O'Hanlon, B., Psiaki, M., Powell, S., Bhatti, J., Humphreys, T. E., Crowley, G., and Bust, G., "CASES: A Smart, Compact GPS Software Receiver for Space Weather Monitoring," *Proceedings of the ION GNSS Meeting*, Institute of Navigation, Portland, Oregon, 2011.

[9] Symmetricom, *Z3801A GPS Receiver User's Guide*, May 2000, `www.leapsecond.com/museum/z3801a/097-z3801-01-iss-1.pdf`.

[10] Schweitzer Engineering Laboratories, *SEL-2401 - Accurate Time Whenever You Need It*, 2005, `http://www.selinc.com/sel-2401`.

[11] Trimble, *Juno Series User Guide*, Oct. 2008, `http://www.trimble.com/junosb_ts.asp`.

[12] Wesson, K., Shepard, D., Bhatti, J., and Humphreys, T. E., "An Evaluation of the Vestigial Signal Defense for Civil GPS Anti-Spoofing," *Proceedings of the ION GNSS Meeting*, Institute of Navigation, Portland, Oregon, 2011.

[13] Humphreys, T. E., Psiaki, M. L., Kitner, P. M., and Ledvina, B. M., "GPS carrier tracking loop performance in the presence of ionospheric scintillations," *Proceedings of the ION GNSS Conference*, Institute of Navigation, Long Beach, California, Sept. 2005.

[14] Kuykendall, P. and Loomis, P. V. W., "In sync with GPS: GPS clocks for the wireless infrastructure," Tech. rep., Trimble Navigation.

[15] Grady, W. M. and Castello, D., "Implementation and application of an independent Texas synchrophasor network," Tech. rep., Schweitzer Eng. Laboratories, Jan. 2010.

[16] Schweitzer, E., Whitehead, D., Guzman, A., Gong, Y., and Donolo, M., "Advanced real-time synchrophasor applications," Tech. rep., Schweitzer Eng. Laboratories, Sept. 2008.