Evaluation of Smart Grid and Civilian UAV Vulnerability to GPS Spoofing Attacks

Daniel P. Shepard, Jahshan A. Bhatti, and Todd E. Humphreys
The University of Texas at Austin
Aaron A. Fansler
Northrop Grumman Information Systems

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.

Jahshan A. Bhatti is pursuing a Ph.D. in the Department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin, where he also received his M.S. and B.S. He is a member of the UT Radionavigation Laboratory. His research interests are in the development of small satellites, software-defined radio applications, space weather, and GNSS security and integrity.

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. He specializes in applying optimal estimation and signal processing techniques to problems in radionavigation. His recent focus is on radionavigation robustness and security.

Aaron A. Fansler serves as cyber critical infrastructure protection (CCIP) program manager for Northrop Grumman Information System. He obtained a Master's degree from Capitol College in information assurance and is currently working on a Ph.D. in information assurance.

ABSTRACT

Test results are presented from over-the-air civil GPS spoofing tests from a non-negligible stand-off distance. These tests were performed at White Sands Missile Range (WSMR) against two systems dependent on civil GPS, a civilian unmanned aerial vehicle (UAV) and a GPS time-reference receiver used in "smart grid" measurement devices. The tests against the civil UAV demonstrated that the UAV could be hijacked by a GPS spoofer by altering the UAV's perceived location. The tests against the time-reference receiver demonstrated the spoofer's capability of precisely controlling timing from a distance, which means a spoofer could manipulate measurements used for smart grid control without requiring physical access to the measurement devices. Implications of spoofing attacks against each of these systems are also given. Recommendations are presented for regulations regarding GPS receivers used in critical infrastructure applications. These recommendations include creating a certification process by which receivers are declared spoof-resistant if they are able to detect or mitigate spoofing attacks in a set of canned scenarios. The recommendations also call for a mandate that only spoof-resistant receivers be used in applications classified by the Department of Homeland Security (DHS) as national critical infrastructure.

I. Introduction

The design of the Global Positioning System came together over Labor Day weekend in 1973. A group of hard-working engineers, mostly Air Force officers, decided over that weekend that the GPS satellites would broadcast two different types of signals, a precise military signal and a so-called clear access or C/A signal. The military signal would later be encrypted to prevent unauthorized use and imitation. But the clear access signal, true to its name, would be freely

Preprint of the 2012 ION GNSS Conference Nashville, TN, September 19–21, 2012

accessible to all. Detailed and accurate specifications for the clear access signal were later distributed to encourage its use.

The early designers of the GPS system, for whose tireless efforts we are all indebted, knew GPS was going to be valuable for civilians across the globe, but they never could have imagined just how valuable. An intentional degradation of the C/A signals called selective availability was discontinued by presidential order in 2000. Instantaneously, every GPS receiver across the globe went from errors the size of a football field to errors the size of a small room. It is hard to overstate the impact of this improvement in accuracy. Before selective availability was turned off, there were no in-car navigation systems giving turn-by-turn directions, because back then civilian GPS could not tell you what block you were on, let alone what street. For geolocation, accuracy matters.

Things have only improved over the last decade. With more ground stations, better algorithms, more open-access signals, and better receivers, civil GPS—the family of open-access signals to which all civilians have access—can now tell you not only what street you are on, but what part of the street. The accuracy, transparency, and low cost of civil GPS have enabled a firestorm of innovation. After 2000, any engineer designing a system for which accurate timing or location was important found GPS to be an almost irresistible option. As a result, civil GPS receivers are built deeply into our national infrastructure: from our smartphones to our cars to the Internet to the power grid to our banking and finance institutions. Some call GPS the invisible utility: it works silently, and for the most part perfectly reliably, in devices all around us of which we are scarcely aware.

However, the same transparency and predictability that has made civil GPS signals so wildly popular has given rise to a significant vulnerability. Transparency and predictability make the civil GPS signals easy to imitate or counterfeit. Civil GPS signals are like Monopoly money: they have a detailed structure but no built-in protection against forgery. The fact that civil GPS is so easy to counterfeit, or "spoof," would not be of importance if GPS were not so popular and its use so widespread. However, this is not the case.

In 2001, the U.S. Department of Transportation (USDOT) evaluated the transportation infrastructure's GPS vulnerability and first raised concern over the threat of GPS spoofers [1]. The USDOT report noted the absence of any off-the-shelf defense against this type of attack 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].

GPS spoofing is the act of producing a falsified version of the GPS signal with the goal of taking control of a target GPS receiver's position-velocity-time (PVT) solution. This is most effectively accomplished when the spoofer has knowledge of the GPS signal as seen by the target receiver so that the spoofer can produce a matched, falsified version of the signal. In the case of military signals, this type of attack is nearly impossible because the military signal is encrypted and therefore unpredictable to a would-be spoofer. The civil GPS signal, on the other hand, is publicly-known and readily predictable.

In recent years, civil GPS spoofing has been recognized as a serious threat to many critical infrastructure applications which rely heavily on the publicly-known civil GPS signal. A number of promising methods are currently being developed to defend against civil GPS spoofing attacks, but it will still take a number of years before these technologies mature and are implemented on a wide scale. Currently, there is a complete absence of any off-the-shelf defense against a GPS spoofing attack.

On invitation from the Department of Homeland Security (DHS), unclassified spoofing tests were performed against two different systems dependent on civil GPS, a civilian unmanned aerial vehicle (UAV) and a GPS time-reference receiver used in "smart grid" measurement devices. These tests took place at White Sands Missile Range (WSMR) on June 19, 2012 during the DHS GYPSY test exercise. In these tests, the capability of a spoofer, developed by the University of Texas at Austin (UT) Radionavigation Lab, to alter the timing and positioning of GPS receivers in these two applications was demonstrated over-the-air from a stand-off distance of about 620 m.

This report details the tests performed at WSMR during the DHS GYPSY test exercise and the spoofer used for the tests. A discussion of the effects of GPS spoofing attacks on the two tested systems is also

provided. Finally, recommendations for regulations on spoofing resistance are presented.

II. Background

A. Civil UAVs

A.1 Iran Drone Incident

In December 2011, Iran captured a U.S. Central Intelligence Agency (CIA) surveillance drone with only minor damage to the undercarriage of the drone, likely due to a rough landing when captured. An Iranian engineer claimed in an interview that "Iran managed to jam the drone's communication links to American operators" causing the drone to shift into an autopilot mode that relies solely on GPS to guide itself back to its home base in Afghanistan. With the drone in this state, the Iranian engineer claimed that "Iran spoofed the drone's GPS system with false coordinates, fooling it into thinking it was close to home and landing into Iran's clutches" [3].

Although the Iranian claims are highly questionable, this incident left many unanswered questions as to the security of GPS systems on unmanned aerial vehicles (UAVs). The CIA drone should have been guiding itself based on the encrypted military GPS signals, which would be incredibly difficult to spoof. However, some experts have conjectured that simultaneous jamming of the military signals and spoofing of the civilian signals might have worked if the drone had been programmed to fall back on the civilian GPS signals in the event that the military signals were jammed. This raises the question: How difficult would it be to spoof a UAV guiding itself based on civilian GPS signals?

A.2 FAA Modernization and Reform Act of 2012

In February 2012, the U.S. Congress passed the FAA Modernization and Reform Act of 2012. According to the Library of Congress summary, this act "requires the Secretary [of Transportation] to develop a plan to accelerate safely the integration by September 30, 2015, of civil unmanned aircraft systems (UASes, or drones) into the national airspace system [and] determine if certain drones may operate safely in the national airspace system before completion of the plan" [4].

Such civilian UAVs would be primarily guided by civil GPS, which has been shown to be readily spoofable in the lab. This would create a significant potential hazard in the national airspace if the problem of civil GPS spoofing is not fixed. Thousands of civilian UAVs (operated by postal services, police departments, research institutions, and others) could populate the skies in only a few years while still being vulnerable to remote hijacking via GPS spoofing. The passing of the FAA Modernization Act further emphasizes the need to examine the vulnerability of UAVs to GPS spoofing.

B. Synchrophasors

As electric power grids continue to expand throughout the world and transmission lines are pushed to their operating limits, the dynamic operation of the power system has become more of a concern and more difficult to accurately model. More effective real-time system control is now seen as key to preventing wide-scale cascading outages like the 2003 Northeast Blackout [5]. For years, electric power control centers have estimated the state of the power system (the positive sequence voltage and phase angle at each network node) from measurements of power flows. But for improved accuracy in the so-called power system state estimates, it will be necessary to feed existing estimators with a richer measurement ensemble or to measure the grid state directly.

Alternating current (AC) quantities have been analyzed for over 100 years using a construct developed by Charles Proteus Steinmetz in 1893, known as a "phasor" [6]. In power systems, the phasor construct has commonly been used for analyzing AC quantities, assuming a constant frequency. A relatively new synchronization technique which allows referencing measured current or voltage phasors to absolute time has been developed and is currently being implemented throughout the world. The measurements produced by this technique are known as "synchronized phasor measurements" or "synchrophasors." Synchrophasors provide a real-time snapshot of current and voltage amplitudes and phases across a power system, and so can give a complete picture of the state of a power system at any instant in time. This makes synchrophasors useful for measurement, analysis, and control of the power grid.

A device used to measure synchrophasors is called

a phasor measurement unit (PMU). In a typical deployment, PMUs are integrated in protective relays and are sampled from widely dispersed locations in the power system network [7]. In order to make accurate measurements of phase angles, PMUs must have a synchronized timing source accurate to better than 26.5 μs according to the IEEE C37.118 Standard "Synchrophasors for Power Systems" [8]. PMUs are synchronized with respect to the common time source of a GPS time-reference receiver to satisfy this accuracy requirement. This raises two questions:

1. Can a civil GPS spoofer cause the time-reference receivers used to synchronize PMUs to violate the IEEE standard for synchrophasor measurements in a realistic scenario?

2. What effects could violating the standard have on control systems reliant on synchrophasor measurements?

III. Civil GPS Spoofing

The spoofer used for these tests was an improved version of the spoofer originally reported in Ref. [2]. A picture of the civil GPS spoofer, developed by the UT Radionavigation Laboratory, is shown in Fig. 1. It is the only spoofer reported in open literature to date that is capable of precisely aligning the spreading codes and navigation data of its counterfeit signals with those of the authentic GPS signals at the target receivers antenna. 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. The 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 downmixes and digitizes GPS L1 and L2 frequencies.

- A DSP board that performs acquisition and tracking of GPS L1 C/A signals, calculates a navigation solution, predicts the L1 C/A databits, and produces a consistent set of up to 14 spoofed GPS L1 C/A signals with a user-controlled 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.

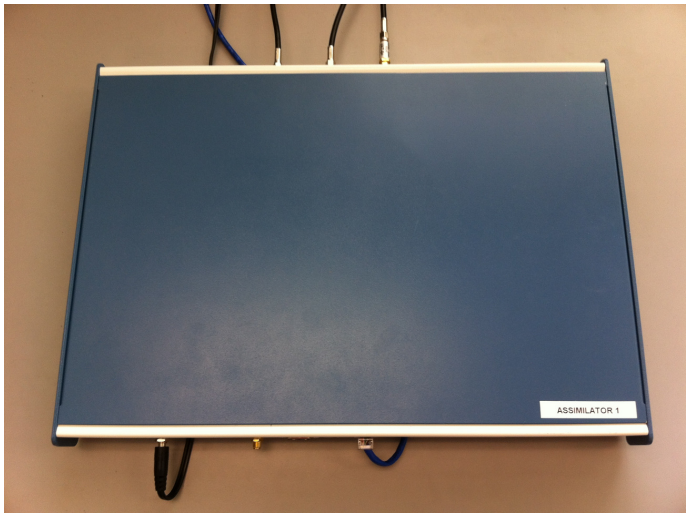

Fig. 1. The Civil GPS Spoofer.

A. Receiver/Spoofer Architecture

The spoofer was designed to operate in conjunction with a software-defined GPS receiver. This design aids the spoofer in producing counterfeit signals which are initially precisely aligned with the authentic signals by leveraging the information obtained about the authentic signals through normal receiver operation. As can be seen from the block diagram of the spoofer in Fig. 2, the spoofer control module utilizes the GPS observables (code phase, carrier phase, and Doppler frequency) and navigation solution output from the coupled receiver. These observables are modified using a linearized measurement model and used to simulate n simulated or "spoofed" GPS signals whose suggested position-velocity-time (PVT) solution is offset, by a user controlled amount, from the navigation solution of the coupled receiver. The spoofer also requires predicted navigation data from the coupled receiver or an external source, which allows the spoofer to produce GPS signals which are nearly indistinguishable from the authentic GPS signals. Additional details on this architecture are provided in Ref. [2] and [9].

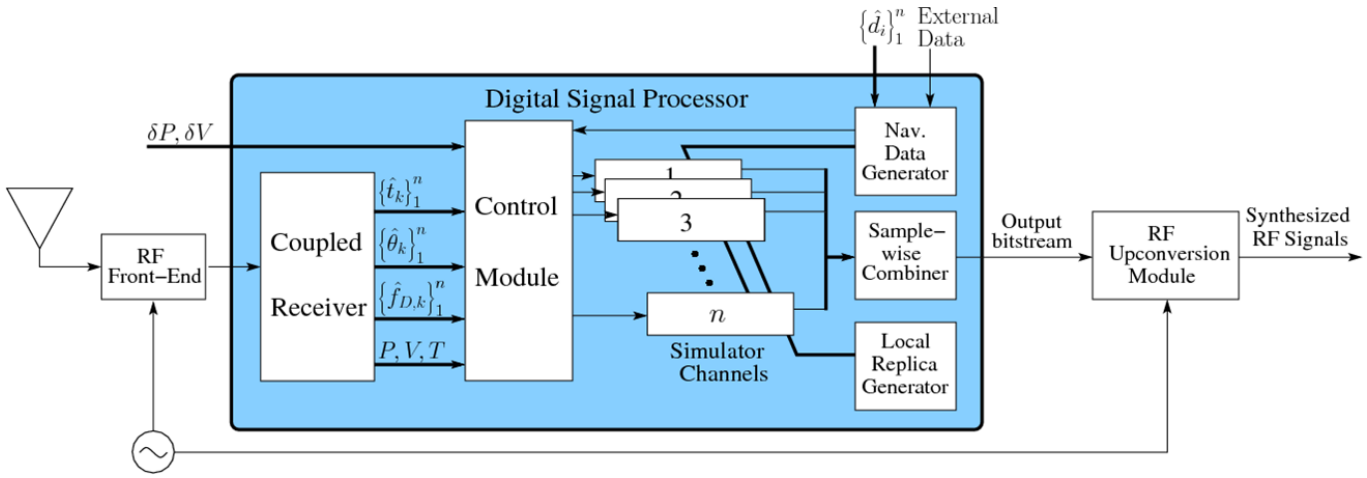

Fig. 2. A block diagram of the Spoofer.

B. Attack Strategy

The spoofer operates by first acquiring and tracking GPS L1 C/A signals to obtain a navigation solution. It then enters its "feedback" mode, in which it produces a counterfeit, data-free feedback GPS signal that is summed with its own antenna input. The feedback signal is tracked 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 navigation data bit library, the spoofer is ready to begin an attack. Initially, it produces signals that are aligned with the authentic signals at the location of the target antenna to within a few meters, but have low enough power that they remain far below the target receiver's 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 can slowly lead the spoofed signals away from the authentic signals, carrying the receiver's tracking loops with it. The target receiver can be considered completely captured when either one of the following are true: (1) each spoofed signal has shifted by 2 μs relative to the authentic signals, or (2) each spoofed signal is at least 10 dB more powerful than the corresponding authentic signal. The latter option ensures that there is no significant interaction between authentic and spoofed signals by simultaneously jamming and spoofing.

The UT spoofer and attack strategy have been tested against a wide variety of civil GPS receivers and have always been successful in commandeering the target receiver. Several of the receivers that have been spoofed are highlighted in Ref. [10].

C. Proximity Spoofing Attack

The spoofing tests performed in the past using the UT spoofer can all be considered to be proximity spoofing attacks. A proximity spoofing attack, as depicted in Fig. 3, is a class of spoofing attacks where the spoofer is located within a few meters of the target receiver, so the distance between the spoofer and target receiver can be neglected. This attack scenario is described in detail in Ref. [2] and significantly decreases the complexity of carrying out an attack. It should be noted that past tests have been performed through-cable or in an RF-shielded enclosure to avoid violating FCC regulations by broadcasting in the GPS band.

D. Spoofing at a Distance

For an attack against a UAV, the only way the spoofer could be assured to be a negligible distance from the target receiver is if the spoofer were attached to the UAV. It is unlikely that this would be the case, so an attack against a UAV will not fall under the category of a proximity spoofing attack. For that matter, physical security of a receiver would often prevent proximity spoofing in most realistic scenarios. This requires the spoofer to consider the effects of spoofing from a non-negligible distance away if precise alignment of the counterfeit and authentic signals is de-

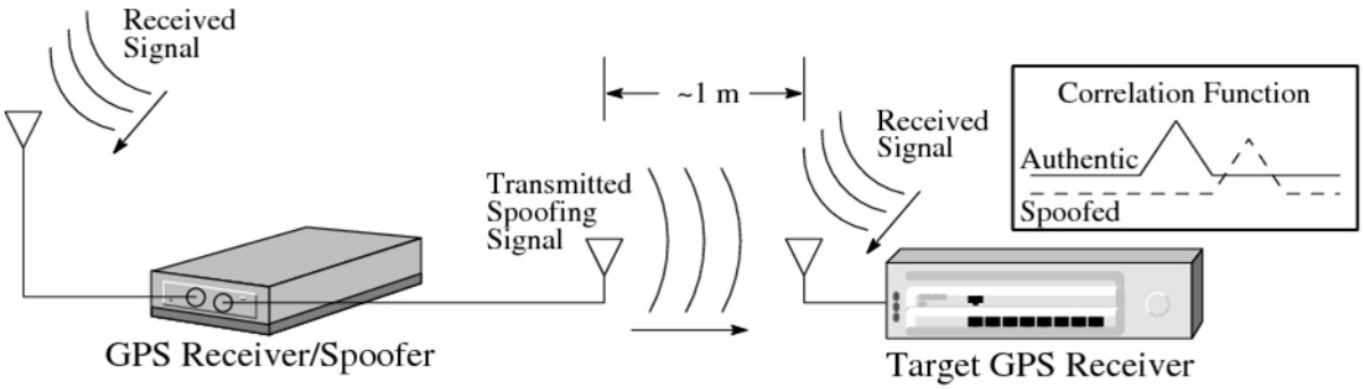

Fig. 3. A diagram of a proximity spoofing attack.

sired. In fact, fine-grained control of a UAV via GPS spoofing is only possible with a meter-level accurate suggested position. Modifications were made to the UT spoofer to account for these effects so that meter-level accurate suggested position was achieved during the tests.

IV. UAV Spoofing Demonstration

A. The UAV

The UAV spoofing tests targeted a UT-owned Hornet Mini UAV supplied by Adaptive Flight, which is shown in Fig. 4. The Hornet Mini is roughly five feet long and weighs about 10 pounds when fully loaded. The Mini's sophisticated avionics package loosely couples an altimeter, a magnetometer, and a MEMS IMU package to a GPS receiver via an extended Kalman filter.

The results of the spoofing tests with the Hornet Mini also apply to other similarly-designed UAVs; those whose navigation systems are centered on civil GPS. The UAVs designed in this way include those used in most non-US-military applications. It should be noted that no special alterations where made to the Hornet Mini for this test—it was in its "as sold" or "stock" configuration.

B. Setup

A schematic of the setup used for the spoofing tests against the civil UAV at WSMR appears in Fig. 5. The spoofer was located on a hilltop with the receive antenna on the far side of the hilltop from the transmit antenna as shown in Fig. 6. The UAV site was

Fig. 4. The Hornet Mini unmanned aerial vehicle (UAV), owned by the UT, used in the spoofing tests.

located in a sandy basin approximately 620 m from the transmit antenna.

C. Procedure

The UAV was commanded by its ground controller to hover approximately 40 feet above ground level at the UAV site. After the initial ground control command was sent, the UAV maintained its hovering position automatically based on the navigation solution of its extended Kalman filter, which is based in part on GPS. At this point in the test procedure, the spoofed signals were not being broadcast: the UAV was only under the influence of the authentic GPS signals.

The spoofer was then commanded to begin transmitting spoofed signals. To ensure seamless capture of the UAV's GPS unit, the code phases of the spoofed signals were aligned to within meters of the authen-

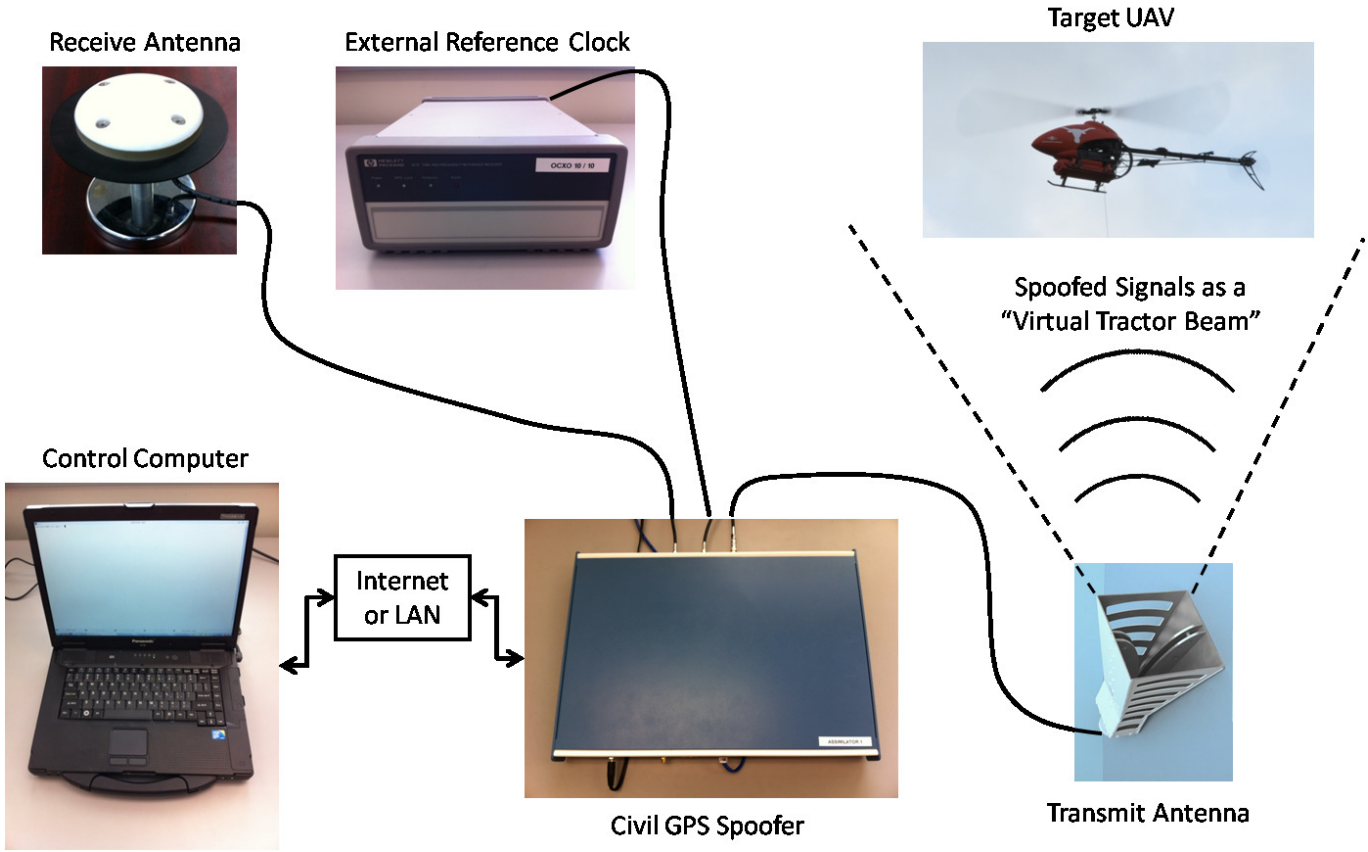

Fig. 5. A schematic of the UAV test setup.

tic signals at the location of the UAV's GPS antenna. The spoofed signals overpowered their authentic counterparts and instantly captured the tracking loops within the UAV's GPS receiver.

Immediately after capture, the spoofer induced a false velocity and corresponding position change in the UAV's GPS receiver, drawing the position reported by the UAV's extended Kalman filter away from the UAV's commanded hover position. To compensate, the UAV's flight controller responded by moving in the opposite direction. A safety pilot was on hand to prevent the UAV from drifting out of control. This was necessary because by commandeering the UAV's GPS receiver, the spoofer operator effectively breaks the UAV autopilot's feedback control loop. The spoofer operator must now act as an operator-in-the-loop, which requires real-time, meter-level knowledge of the UAV's true location.

D. Results

Between tests at WSMR and UT, the spoofer demonstrated short-term 3-dimensional control of the UAV. Thus, it is possible to hijack a civil UAV—in this case, a fairly sophisticated one—by civil GPS spoofing.

Interestingly, the Hornet Mini relies only on its altimeter for direct measurements of its vertical position; the GPS-measured vertical position is ignored. This can be done with reasonable accuracy because of the Hornet Mini's short flight endurance (about 20 minutes). However, the GPS vertical velocity does affect the extended Kalman filter's vertical coordinate estimate because the filter propagates GPS velocity measurements through a UAV dynamics model to form an a priori vertical estimate that gets updated with the altimeter measurements. This dependence on GPS velocity allowed the spoofer operator to force the UAV vertically downward in dramatic fashion in the final three capture demonstrations.

E. Implications

These tests have demonstrated that civilian UAVs will be vulnerable to control by malefactors with a civil GPS spoofer looking to hijack or crash these UAVs unless their vulnerability to GPS spoofing is

Fig. 6. Aerial view of the test site showing the spoofer location on a hilltop and the UAV site approximately 0.62 kilometers away.

addressed. There are several reasons why someone may want to spoof a drone including fear over drones invading people's privacy. This poses a significant safety concern that could result in mid-air collisions with other aerial vehicles or buildings, not to mention loss of property.

Constructing from scratch a sophisticated GPS spoofer like the one developed by UT is not easy, nor is it within the capability of the average anonymous hacker. It is orders of magnitude harder than developing a GNSS jammer. Nonetheless, the trend toward software-defined GNSS receivers for research and development, where receiver functionality is defined entirely in software downstream of the A/D converter, has significantly lowered the bar to developing a spoofer in recent years. As a point of reference, we estimate that there are more than 100 researchers in universities around the world who are well-enough versed in software-defined GPS that they could develop a sophisticated spoofer from scratch with a year of dedicated effort.

More worrisome is the fact that one does not have to build a sophisticated spoofer like ours, capable of aligning its signals precisely with authentic signals at the location of a chosen target, to spoof a civil GPS receiver. A low-cost off-the-shelf GPS signal simulator would not permit the kind of seamless attack we carried out, but would be adequate to confuse and disrupt the navigation system of a commercial UAV.

V. GPS Time-Reference Receiver Spoofing Demonstration

A. Prior Tests

In December 2011, the University of Texas at Austin and Northrop Grumman Information Systems performed laboratory spoofing tests against a GPS time-reference receiver supplying timing to a PMU. The minimum threshold for success in these spoofing tests was to show that a GPS spoofer could force a PMU to violate the IEEE C37.118 Standard "Synchrophasors for Power Systems" [8]. The standard requires a phase angle error of less than 0.573°, which can be equivalently and indistinguishably caused by a timing error of 26.5 μs.

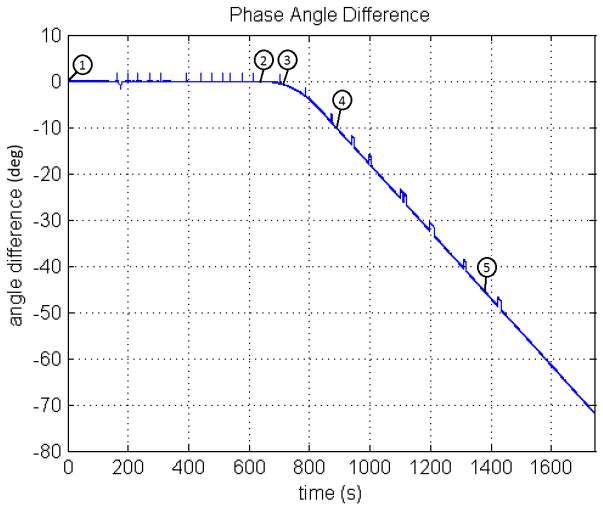

Fig. 7. A plot of the phase angle difference between the reference and spoofed PMUs. Normally the phase angle difference would be nearly zero in the absence of a spoofing attack. Point 1 marks the start of the test. Point 2 marks the point at which the spoofer has completely captured the target receiver. Point 3 marks the point at which the IEEE C37.118 Standard has been broken. Point 4 marks the point at which the spoofer-induced velocity has reached its maximum value for the test. Point 5 marks the point at which the spoofed signal was removed.

In these tests, the phase angle of the spoofed PMU was monitored as well as the phase angle from a non-spoofed PMU in the same room. Figure 7 shows the measured phase angle difference between the reference PMU, which was fed the true GPS signal, and the spoofed PMU throughout one entire test. This value would normally be less than a few degrees in the absence of spoofing, since the two PMUs are co-located. After the initial ten minute capture and carry-off, which proceeds slowly to avoid detection, the spoofer accelerates its timing carry-off and the reference and spoofed phase angles quickly diverge.

Figures 8 through 12 show pictures of an oscilloscope and the synchrophasor screen at different times throughout the test. The oscilloscope shows two pulse-per-second (PPS) signals, with the upper yellow pulse coming from a reference clock being fed true GPS and the lower blue pulse coming from the spoofed timing receiver. Both PPS signals are initially aligned with each other, as seen in 8. The synchrophasor screen displays the PMU phase angle data in real-time as phasors with the nominal 60 Hz operating frequency subtracted from the phase angle. The red and green phasors show the phase data from the reference and spoofed PMUs respectively. These phasors are within a few degrees of each other at the beginning of the test, as seen in 8.

At the time shown in Fig. 10, the IEEE C37.118 Standard was broken. The spoofer was easily able to break this standard and go much further. The spoofer-induced phase angle error exceeded 10^o within 15 minutes of the start of the test, as shown in Fig. 11. By the end of the test, the spoofer-induced phase angle error exceeded 70^o, as shown in Fig. 7.

This test demonstrated that a proximity spoofing attack against a PMU can induce large, spoofer-controlled errors in the phase angle measured by the PMU in a relatively short period of time without causing any alarms in the system. A complete description of these tests and their implications can be found in Ref. [11].

B. Setup

The setup for the WSMR time-reference receiver spoofing test was exactly the same as for the UAV spoofing tests, shown in Fig. 5, on the spoofer end, and the target site was also at the same location, shown in Fig. 6. At the target site, there were two GPS time-reference receivers. The first time-reference receiver was representative of the ones used for PMU networks and served as the target of the spoofing attack. The other time-reference receiver was used as a time reference during the testing by unplugging the GPS antenna before the spoofing attack began. This forced the receiver into its "holdover" or GPS-denied mode. While in holdover mode, the time-reference receiver was able to ride through the spoofing attack using its highly stable ovenized crystal oscillator (OCXO) to maintain accurate timing.

C. Procedure

Before the spoofing attack began, the time alignment of the two time-reference receivers was observed on an oscilloscope using the IRIG-B output from the target receiver and the PPS output from the reference receiver. The oscilloscope was set to trigger on the PPS output from the reference receiver. Once the two receivers agreed to within 100 ns, which is typical for these two receivers, the reference receiver was unplugged from the antenna and allowed to transition into holdover mode. Data was recorded from the oscilloscope to demonstrate this time alignment.

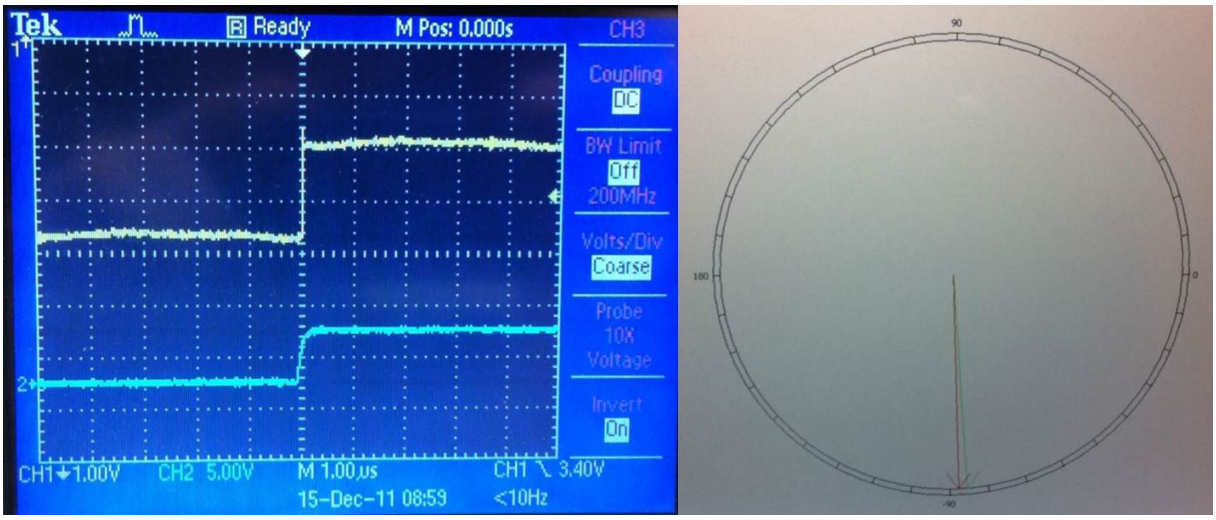

Fig. 8. Pictures of the oscilloscope (left) and synchrophasor (right) screen at the start of the test, which is marked as point 1 in Fig. 7.

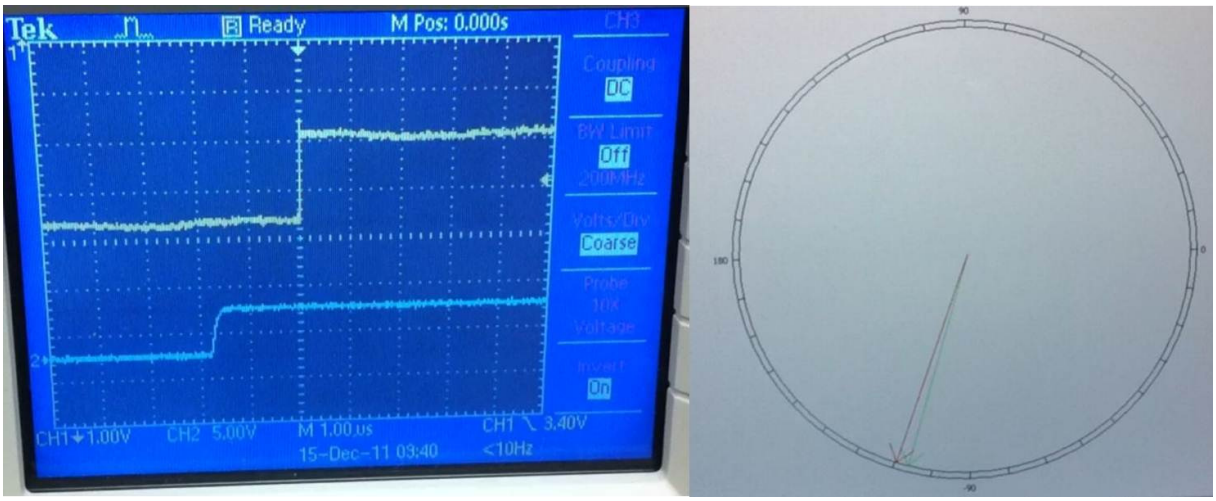

Fig. 9. Pictures of the oscilloscope (left) and synchrophasor (right) screen at about 620 seconds into the test, which is marked as point 2 in Fig. 7.

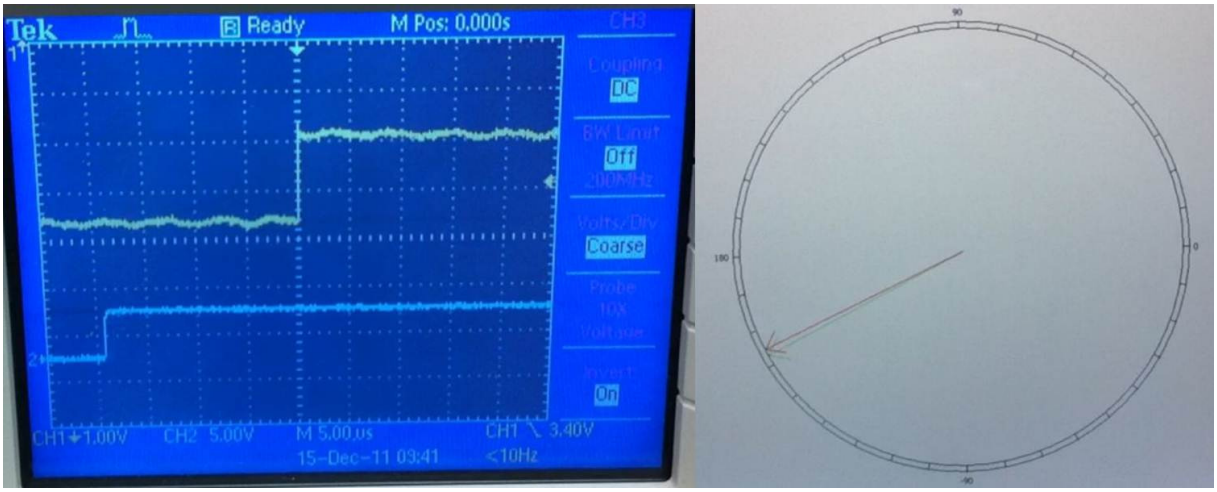

Fig. 10. Pictures of the oscilloscope (left) and synchrophasor (right) screen at about 680 seconds into the test, which is marked as point 3 in Fig. 7.

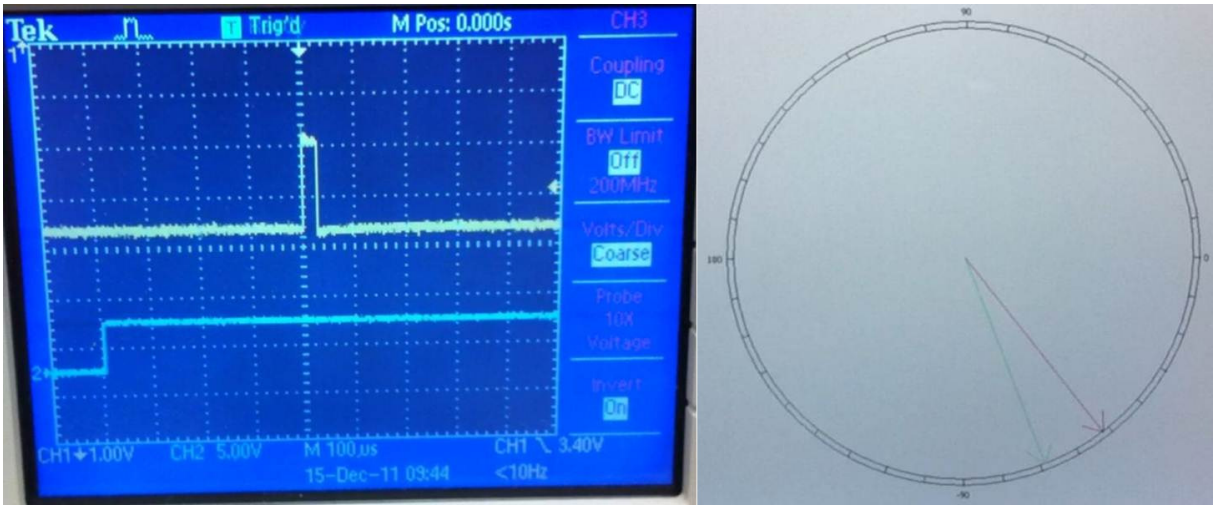

Fig. 11. Pictures of the oscilloscope (left) and synchrophasor (right) screen at about 870 seconds into the test, which is marked as point 4 in Fig. 7.

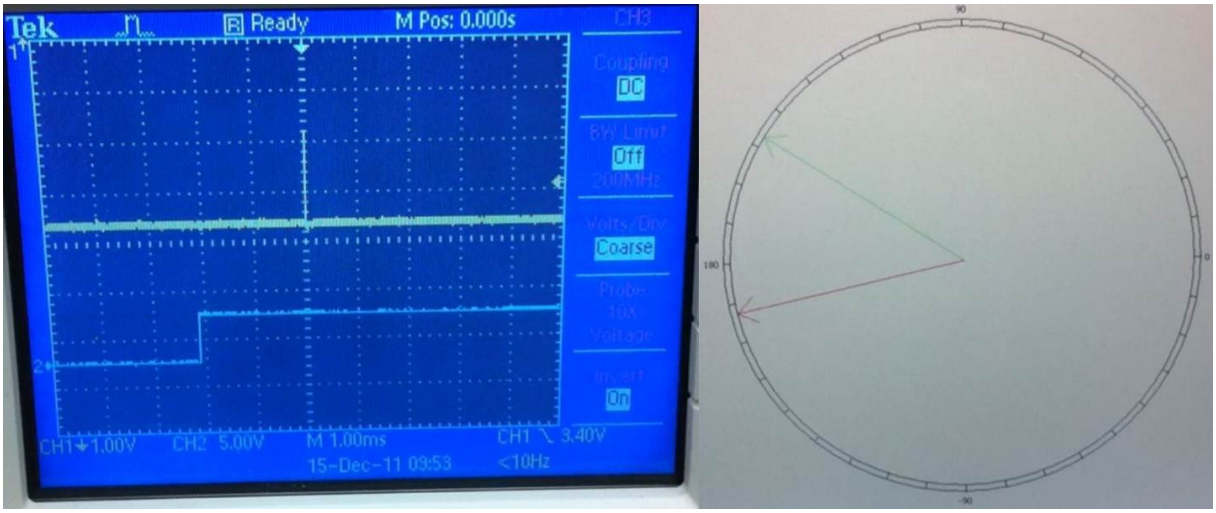

Fig. 12. Pictures of the oscilloscope (left) and synchrophasor (right) screen at about 1370 seconds into the test, which is marked as point 5 in Fig. 7.

At this point, the spoofer began transmitting spoofed signals that were initially nearly perfectly aligned with the authentic signals at the target site. The spoofed signals overpowered their authentic counterparts and instantly captured the tracking loops within the target receiver. The spoofer then began to drag the timing of the target receiver away from the truth until it reached 1 μs of induced timing error. This was chosen to demonstrate that the spoofer had precise control over the target receiver's timing. Data was recorded from the oscilloscope to show that a 1 μs induced timing error was achieved.

Finally, the spoofer was commanded to cease transmitting the spoofed signals. Once the target receiver reacquired the authentic signals and corrected its timing, data was recorded from the oscilloscope to demonstrate that the reference receiver did not drift significantly in timing during the test.

D. Results

Figure 13 shows the data taken from the oscilloscope from before the spoofing attack began. This demonstrates that the two time reference receivers agree to within 100 ns nominally. Figure 14 shows the data taken from the oscilloscope from the end of the spoofing test, where the spoofed time-reference receiver has a spoofer-induced timing error of almost exactly 1 μs. This shows that the spoofer was able to precisely control the timing of the spoofed receiver during the test.

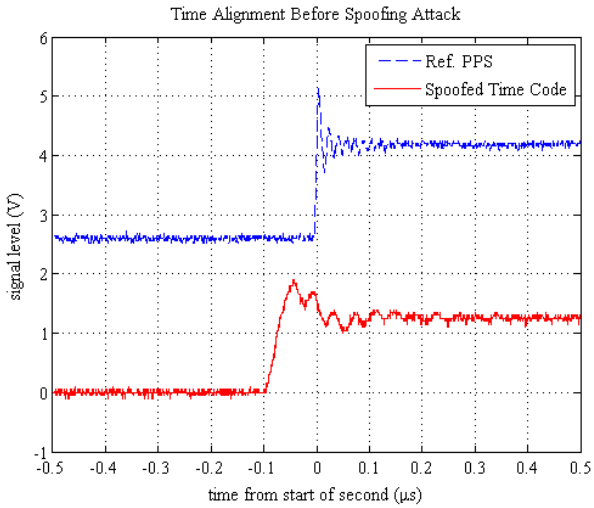

Fig. 13. Time alignment of the reference PPS (top blue dashed line) and the spoofed IRIG-B time code (bottom red line) before the spoofing attack began.

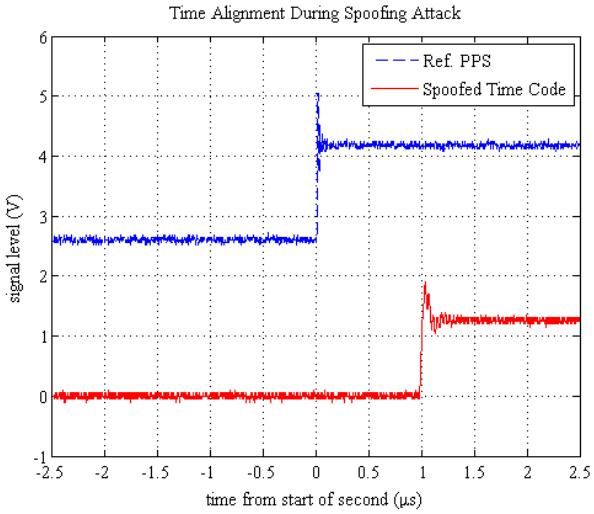

Fig. 14. Time alignment of the reference PPS (top blue dashed line) and the spoofed IRIG-B time code (bottom red line) at the end of the spoofing attack.

Figure 15 shows the data taken from the oscilloscope from after the spoofing test, once the spoofed receiver reacquired the authentic signals and corrected its timing. This demonstrates that the reference receiver did not drift significantly in timing during the test, which means that any change in relative timing between the reference and spoofed receivers can be attributed to the effects of the spoofer.

E. Implications

In a practical scenario, a malefactor may seek to subvert the control objectives of electric power authori-

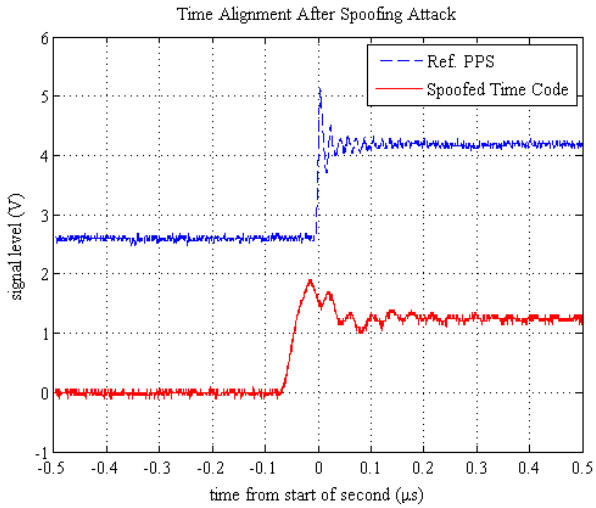

Fig. 15. Time alignment of the reference PPS (top blue dashed line) and the spoofed IRIG-B time code (bottom red line) after the spoofing attack ended.

ties by altering their perception of the current state of the power grid. The end goal of the malefactor may be to cause damage to power grid equipment or local blackouts. Between this demonstration of timing control from a distance and the prior tests described in detail in Ref. [11], it has been demonstrated that a sophisticated spoofing attack can alter the phase angle measurements of a PMU network without needing physical access to the devices themselves. The simplest synchrophasor-based control scheme relies solely on phase angle differences between two PMUs as an indicator of a fault condition. Thus, a malefactor could accomplish his goals by targeting important power grid nodes (i.e. areas with high power flow) with a GPS spoofing attack which alters the timing in a way that increases the phase angle differences between nodes in the area. This type of attack would likely be indistinguishable from an actual fault and cause corrective actions to be taken when none are necessary.

PMUs are not currently being used for control purposes in the U.S., but the industry and government are pushing for more efficient distribution of power which will require the accuracy and data rates that PMUs provide for state estimation of the power grid. However, other countries are already beginning to implement synchrophasor-based control schemes. One example of a currently operational synchrophasor-based control system is the Chicoasen-Angostura transmission link in Mexico [12]. This transmission line links large hydroelectric generators in Angos-

tura to large loads in Chicoasen through two 400-kV transmission lines and one 115-kV transmission line. PMUs are stationed at each end of the transmission line and are setup to automatically trip the hydroelectric generators offline in the event that the phase angle difference between the two PMUs exceeds 10^o. This system was implemented to protect the generators against fault conditions. If a spoofer were to attack this system in Mexico or a similar implementation elsewhere, then the spoofer could easily cause an unnecessary generator trip in a matter of minutes.

Beyond tripping a single generator, there is potential for the effects of a spoofing attack to propagate through the grid and cause cascading faults across the grid. This was best demonstrated by the 2003 Northeast Blackout, which originated with the tripping of a single transmission line [5]. In a little more than an hour, this event cascaded into a large scale blackout that left 50 million people without power for four days and cost an estimated six billion dollars. Although future control systems are being designed to prevent an event from scaling to this magnitude, a single spoofer targeting the right node would likely still have wide reaching effects if a malefactor had knowledge of the power grid architecture. Additionally, a network of spoofers carrying out a coordinated spoofing attack against various nodes on the power grid could greatly increase the area of effect.

VI. Fixing the Problem of GPS Spoofing

There is no quick, easy, and cheap fix for the civil GPS spoofing problem. Moreover, not even the most effective GPS spoofing defenses are foolproof. In contrast to message authentication, such as is used to sign data transmitted across the Internet, the security of GPS signal authentication is much weaker and demands a probabilistic model. Nonetheless, there are many possible remedies to the spoofing problem that, while not foolproof, would vastly improve civil GPS security. These defenses include placing cryptographic signatures in the navigation messages or spread-spectrum codes on either the wide-area augmentation system (WAAS) or GPS satellites, antenna-based defenses, and jamming detectors. A discussion of the advantages and disadvantages of some of these defences is given in Ref. [13]. The ideal spoofing defense is one which:

1. would reliably detect a sophisticated spoofing attack, such as the one conducted at WSMR, with a low probability of false alarm

2. could be implemented in the short term

3. would not significantly increase the cost of a GPS-based navigation system

4. would be applicable to a broad range of GPS dependent systems

VII. Recommendations

It is the authors' recommendation that for non-recreational operation in the national airspace, civil UAVs exceeding 18 lbs be required to employ navigation systems that are spoof-resistant. Additionally, the authors recommend that GPS-based timing or navigation systems having a non-trivial role in systems designated by DHS as national critical infrastructure be required to be spoof-resistant.

Resistance to spoofing will be defined through a series of canned attack scenarios that can be recreated in a laboratory setting [14]. A navigation system is declared spoof-resistant if, for each attack scenario, the system is either unaffected by or able to detect the spoofing attack. Spoofing detection combined with an appropriate GPS-denied mode for the UAV to fall back on will significantly increase the difficulty of mounting a successful spoofing attack against a UAV. Timing receivers could use a spoofing detection mechanism to force themselves into a holdover mode that relies on its local oscillator, like the receiver used as a reference in the timing tests, and send an alert that a spoofing attack is occurring.

Finally, the authors recommend that a cryptographic authentication signature be developed and implementated for one of the existing or forthcoming civil GPS signals. The signature should at minimum take the form of a digital signature interleaved into the navigation message stream of the WAAS signals. A better plan would be to interleave the signature into the CNAV or CNAV2 GPS navigation message stream like the signature described in Ref. [15]. The best plan for implementing a cryptographic authentication signature would be to implement the signature as an spread-spectrum security code (SSSC) interleaved into the spreading code of the L1C data channel like the signature described in Ref. [16]. Inclusion of a cryptographic signature would greatly aid manufac-

turers in developing receivers that are spoof-resistant.

VIII. Conclusions

Test results presented herein demonstrate that a GPS spoofer can alter a civil UAV's perception of its location and a time-reference receiver's perception of the current time from an appreciable distance away. The GPS receivers in both of these tests reported no alarms during the tests to indicate that they suspected their position-velocity-time (PVT) solution was anything other than nominal.

It was demonstrated that a civil UAV could be "steered" by a spoofer by moving its perceived location in the opposite direction of the desired motion. Coarse, short-term control of the UAV was demonstrated in all directions (east, north, and up) during the tests. Since the spoofer did not have real-time feedback of the UAV's current position and velocity, long-term control was unachievable during these tests. However, a medium-sized radar system could be used to provide this feedback, and a control loop could be designed within the spoofer to provide stable control of the UAV. With the passage of the FAA Modernization Act of 2012, civil UAVs could occupy the national airspace within the decade. If the issue of civil GPS spoofing is not fixed before then, then civil UAVs would pose a significant safety concern in the national airspace that could result in mid-air collisions with other aerial vehicles or buildings, not to mention loss of property.

One critical infrastructure application that will soon use GPS time-reference receivers is the power grid. PMUs use time-reference receivers to time stamp their measurements, which allows power grid operators to get a snapshot of the current state of the grid including phase angles. PMUs are a technology that will revolutionize power grid control and pave the way for more efficient power distribution. However, it has been demonstrated in Ref. [11] that a spoofing attack can induce arbitrarily large errors in the PMU-measured phase angles by inducing timing errors in the time-reference receiver driving the PMU. This fact combined with the demonstrations of spoofing from a distance presented herein proves feasibility of a spoofing attack against a PMU in which the spoofer does not require close proximity to the PMU. Altering of PMU-measured phase angles could cause power grid control systems to unnecessarily trip generators or transmission lines. These effects would likely cause local area blackouts and have the potential for causing damage to power grid equipment. There also exists the potential for the effects to cascade into large scale blackouts similar to the 2003 Northeast Blackout.

There is no quick, easy, and cheap fix for the civil GPS spoofing problem. However, many promising techniques that, while not foolproof, would vastly improve civil GPS security have been and are being developed. These defenses include placing cryptographic signatures in the navigation messages or spread-spectrum codes on either the WAAS or GPS satellites, antenna-based defenses, and jamming detectors.

It is the authors' recommendation that for non-recreational operation in the national airspace, civil UAVs exceeding 18 lbs be required to employ navigation systems that are spoof-resistant. Additionally, the authors recommend that GPS-based timing or navigation systems having a non-trivial role in systems designated by DHS as national critical infrastructure be required to be spoof-resistant. Resistance to spoofing will be defined through a series of standardized tests that require the receiver to detect or mitigate the spoofing attack. This combined with regulations concerning GPS-denied modes for systems reliant on GPS would greatly increase the difficulty of mounting a successful spoofing attack. Finally, the authors recommend that a cryptographic authentication signature be developed and implementated for one of the existing or forthcoming civil GPS signals. Inclusion of a cryptographic signature would greatly aid manufacturers in developing receivers that are spoof-resistant.

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] Rawnsley, A., "Iran's Alleged Drone Hack: Tough, but Possible," Dec. 2011, `http://www.wired.com/dangerroom/2011/12/iran-drone-hack-gps/#`.

[4] "Bill Summary & Status 112th Congress (2011 - 2012) H.R.658 CRS Summary," Feb. 2012, `http://thomas.loc.gov/cgi-bin/bdquery/z?d112:HR00658:@@@D&summ2=m&`.

[5] "Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommenda-

tions," Tech. rep., U.S.-Canada Power System Outage Task Force, April 2004.

[6] "Charles P. Steinmetz," http://www.britannica.com/EBchecked/topic/565056/Charles-Proteus-Steinmetz.

[7] Phadke, A. G. and Thorp, J. S., editors, *Synchronized Phasor Measurements and Their Applications*, Springer, New York, 2008.

[8] "IEEE Standard for Synchrophasors for Power Systems," 2005, IEEE Std. C37.118 Revision 1344–1995.

[9] Humphreys, T. E., Bhatti, J., and Ledvina, B., "The GPS Assimilator: a Method for Upgrading Existing GPS User Equipment to Improve Accuracy, Robustness, and Resistance to Spoofing," *Proceedings of the ION GNSS Meeting*, Institute of Navigation, Portland, Oregon, 2010.

[10] Shepard, D. and Humphreys, T. E., "Characterization of Receiver Response to a Spoofing Attack," *Proceedings of the ION GNSS Meeting*, Institute of Navigation, Portland, Oregon, 2011.

[11] Shepard, D. P., Humphreys, T. E., and Fansler, A. A., "Evaulation of the Vulnerability of Phasor Measurement Units to GPS Spoofing," *International Journal of Critical Infrastructure Protection*, Dec. 2012, to appear.

[12] Schweitzer, E. O., Guzman, A., Altuve, H. J., and Tziouvaras, D. A., "Real-Time Synchrophasor Applications for Wide-Area Protection, Control, and Monitoring," Tech. rep., Schweitzer Eng. Laboratories, 2009.

[13] Humphreys, T. E., "Statement on the vulnerability of civil unmanned aerial vehicles and other systems to civil GPS spoofing," `http://homeland.house.gov/sites/homeland.house.gov/files/Testimony-Humphreys.pdf`, July 2012.

[14] Humphreys, T. E., Shepard, D., Bhatti, J., and Wesson, K., "A Testbed for Developing and Evaluating GNSS Signal Authentication Techniques," 2012, in preparation; available at `http://radionavlab.ae.utexas.edu/testbed`.

[15] Wesson, K., Rothlisberger, M., and Humphreys, T. E., "Practical Cryptographic Civil GPS Signal Authentication," *NAVIGATION, Journal of the Institute of Navigation*, Vol. 59, No. 3, 2012, pp. 177–193.

[16] 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.