

Evaluation of the Vulnerability of Phasor Measurement Units to GPS Spoofing Attacks

Daniel P. Shepard (*dshopard.ut@gmail.com*) and Todd E. Humphreys
(*todd.humphreys@mail.utexas.edu*)
The University of Texas at Austin
Aaron A. Fansler (*aaron.fansler@ngc.com*)
Northrop Grumman Information Systems

ABSTRACT

Test results are presented from GPS spoofing tests against Phasor Measurement Units (PMUs) to demonstrate their vulnerability to spoofing attacks. A GPS spoofer can manipulate the timing of a PMU by broadcasting a falsified GPS signal and forcing the time reference receiver that is providing timing for the PMU to track the falsified signal. This spoofer-induced timing offset creates a corresponding change in the phase angle measured by the PMU.

A particular synchrophasor-based automatic control scheme currently implemented in Mexico is described. It is shown that a generator trip could be falsely activated by a GPS spoofing attack in this system, thus highlighting the threat of spoofing a PMU. A description of the events that led to the 2003 northeast blackout is provided as an example of a potential worst case scenario where the legitimate or false tripping of a single generator or transmission line could lead to cascading faults and a large scale blackout.

I. Introduction

The generation, transmission, and distribution of electric power make the power grid the most critical of critical infrastructure in the United States. Past real-world events and numerous government demonstrations have shown just how vulnerable the electric power infrastructure can be, not only to natural disasters, but more importantly to malicious cyber activity which is on the rise. In the past, the consequences of power disruption were annoyance and some economic cost; future disruptions resulting from intentional malicious activity could cascade into crippling failures.

Effective operation of a country's energy infrastructure (electric power, oil, and natural gas production, transmission, and distribution) is critical to the health and safety, national security, and economic viability of that nation. Cyber threats have become more of a concern because they now rival the consequences of physical attacks.

Every element of critical infrastructure, including the electric power industry, originally operated without an external time reference. However, over the past decade, industry has seen an explosion in the use of accurate, synchronized time incorporated into their controlling networks. In the electric power infrastructure, accurate timing signals are being exploited in systems from the generation plant down to the distribution substation and now down to individual smart grid component.

The value of time synchronization is best understood by recognizing that the power grid is a single, complex, interconnected and interdependent network. What happens in one part of the grid affects operation elsewhere. Effects will also be felt in other systems reliant on stable power, much like what was observed in the 2003 Northeast Blackout [1].

With the transition to smart grid technologies and a unified synchronized "grid," the potential for catastrophic cascading failures increases if proper control measures are not implemented. Time-synchronized measurements are changing the way

electric power systems are controlled to protect against these events. Phasor measurement units (PMUs) have recently emerged as one technology which has the potential to one day anticipate failures, making it possible to take remedial actions before failures spread across the network [2].

PMUs rely on GPS to provide accurate, synchronized time across the power grid. This reliance of PMUs on GPS for time synchronization creates a vulnerability to a particular type of malicious attack called GPS spoofing [3]. 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 [4]. 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 [3].

Northrop Grumman Information Systems (NGIS) and the University of Texas (UT) conducted a functional test and evaluation (FT&E) of the effects a spoofed GPS timing signal would have on synchrophasors. GPS spoofing attacks were performed, both through cable and over-the-air inside an RF shielded tent, against a GPS time reference receiver which provided timing for a PMU. The goal was to determine if adverse effects could be produced on a sensitive timing-signal-dependent network such as a Supervisor Control and Data Acquisition (SCADA) network and the network devices such as PMUs.

The minimum threshold for success was to show that a GPS spoofer could force a PMU to violate the IEEE C37.118 Standard "Synchrophasors for Power Systems" [5]. The Standard defines accuracy as a vectorial difference between the measured and expected value of the phasor for the measurement at a given instant of time, called the total vector error (TVE). TVE blends together

three possible sources of error: magnitude, phase angle, and timing. An error in timing appears identical to an error in phase angle. Without timing and magnitude errors, a phase angle error of 0.573° corresponds to a 1% TVE, which is the maximum allowable TVE by the IEEE C37.118 Standard [6]. This phase angle error could be equivalently and indistinguishably caused by a timing error of $26.5 \mu\text{s}$, which was chosen as the threshold for success in the spoofing tests.

II. Background

A. Synchrophasors

As electric power grids continue to expand throughout the world and as 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 a key to preventing wide-scale cascading outages like the 2003 Northeast Blackout [1].

For years, electric power control centers have estimated the state of the power system (the positive sequence voltage and 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” [7]. 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 control, measurement, and analysis of the power system.

In a typical deployment, synchrophasors are integrated in protective relays and are sampled from widely dispersed locations in the power system network. They are synchronized with respect to the common time source of a global positioning system (GPS) clock. Synchrophasors are basically measurements of AC voltage (or current) and absolute phase angle, made at any selected point in an electric transmission or distribution system [2].

B. GPS Spoofing

GPS spoofing is the act of producing a falsified version of the GPS signal with the goal of taking control of a 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. On the other hand, the civil GPS signal is publicly-known and readily predictable.

In recent years, civil GPS spoofing is becoming 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.

III. The Spoofer

The civil GPS spoofer used for these tests, shown in Fig. 1, is an advanced version of the spoofer reported in [3]. It is the only spoofer reported in




Fig. 1. The Civil GPS Spoofing.

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. 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 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 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.

The spoofer works by first acquiring and tracking GPS L1 C/A and L2C signals to obtain a navigation solution. It then enters its “feedback” mode, in which it produces a counterfeit, data-free feed-

back 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. It produces signals that are initially nearly perfectly aligned with the authentic signals but with low enough power that they remain far below the victim 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 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 $2 \mu\text{s}$ in time away from the authentic signals, the victim receiver has been completely captured.

The spoofer and attack strategy have been tested against a wide variety of GPS receivers and has always been successful in spoofing the target receiver. Several of the receivers that have been spoofed are highlighted in Ref. [8].

IV. Test Setup

Figure 2 shows a schematic of the setup used for the open-air tests. The signals received at the roof were routed into the spoofer for use in producing the counterfeit signals and into the RF shielded tent for rebroadcasting. The counterfeit signals were also routed into the tent for broadcasting. In addition to the antennas broadcasting the authentic and counterfeit signals, a third antenna was setup inside the tent to receive the combination of authentic and spoofed signals. This setup is representative of an actual attack scenario where the malefactor does not have physical access to the victim receiver’s antenna input but rather broadcasts the spoofed signals over-the-air. Figure 3

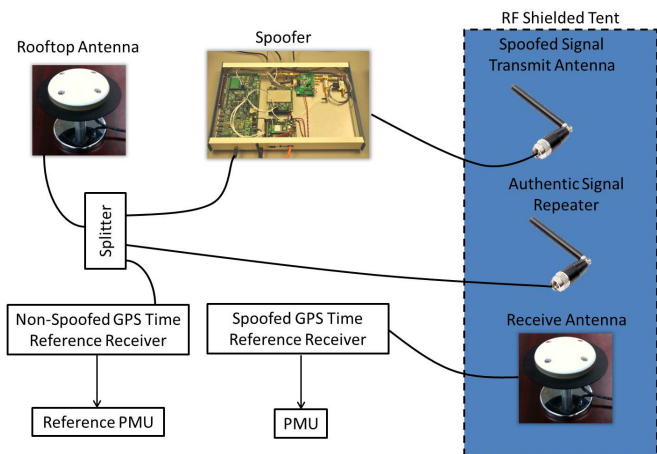


Fig. 2. Schematic of the test setup.



Fig. 4. Picture of the transmit antennas inside the RF shielded tent with one repeating the authentic signal and the other broadcasting the spoofed signal.



Fig. 3. Picture of the outside of the RF shielded tent with cables for the antennas.



Fig. 5. Picture of the receive antenna inside the RF shielded tent which was pulling in both the authentic and spoofed signals to feed to the victim receiver.

shows the outside of the tent where the cables for the transmit and receive antennas were being fed into the tent. Figures 4 and 5 show the transmit and receive antennas respectively as they were set up at opposite ends inside the tent. For cable-only tests, the entire setup inside the tent was replaced with a signal combiner that summed the authentic and spoofed signals.

The combined authentic and spoofed signals were fed to the victim GPS time reference receiver. The output timing signal from the victim receiver was used as the synchronization reference for one PMU, whereas a second PMU was given timing from a separate GPS time reference receiver that was tracking only authentic GPS signals. Since the PMUs were in the same room and measured the local voltage and carrier phasors, both PMUs would report roughly the same phasor measurements under normal circumstances. Thus, any

significant differences in the phase angle measurements between the two PMUs could be attributed to the effects of spoofing.

V. Test Results

Both the cable-only and the over-the-air spoofing attacks were successful in leading the PMU phase measurements off from the truth. Figure 6 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

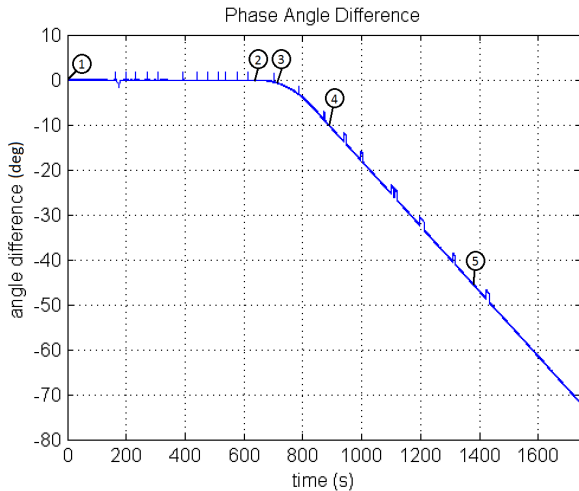


Fig. 6. A plot of the phase angle difference between the reference and the spoofed PMUs. Normally the phase angle difference would be nearly zero in the absence of a spoofing attack. Point 1 marks the start of the test. Point 2 marks the point at which the spoofer has completely captured the victim receiver. 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.

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 carry-off and the reference and spoofed phase angles quickly diverge.

Figure 7 shows pictures of an oscilloscope and the Synchronwave screen at the start of 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. The Synchronwave 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.

Figure 8 shows pictures of the Oscilloscope and the Synchronwave screen at about 620 seconds into

the test. At this point, the spoofer has moved the victim receiver $2 \mu\text{s}$ off in time and has completely captured the receiver. The delicate initial capture-and-carry-off is performed at a slow rate to suppress any evidence of the spoofer's presence. However, this process could be done quicker because the receiver was not looking for such evidence of foul play. At this stage of the test, there is not yet any significant difference between the two phasors on the Synchronwave screen, since the spoofed time offset remains relatively small. The oscilloscope, however, reveals that the PPS output from the victim receiver has moved by about $2 \mu\text{s}$ relative to the reference PPS. At this point, the spoofer begins to accelerate the victim receiver's time solution at a distance-equivalent rate of 4 m/s^2 until it reaches a final distance-equivalent velocity of 1000 m/s . Distance-equivalent velocity can be converted into the actual time rate of change of time by dividing by the speed of light.

The acceleration segment of the attack must be tailored to the individual receiver's ability to track the spoofer-induced dynamics [8]. Otherwise, the spoofer risks losing control of the victim receiver's tracking loops by moving too quickly for the receiver to track or raising alarms. Alternatively, a malefactor could survey possible GPS time reference receiver's that might be used and tailor the spoofing attack such that any of the receivers would track and believe the spoofed signals. This would place severe limits on the spoofer's ability to manipulate timing, but would not make the attack impossible or implausible.

Figure 9 shows pictures of the oscilloscope and the Synchronwave screen at about 680 seconds into the test. At this point, the spoofer has broken the IEEE C37.118 Standard for PMUs, which requires accuracy in the measured phase angle of 0.573° [6]. This demonstrates a significant vulnerability for PMU-based monitoring and control, since these applications leverage the accuracy supposedly guaranteed by the standard. There is yet no noticeable difference on the Synchronwave screen, but the oscilloscope clearly shows that the victim receiver has now been offset in time by about $20 \mu\text{s}$.




Fig. 7. Pictures of the Oscilloscope (left) and Synchrowave (right) screen at the start of the test, which is marked as point 1 in Fig. 6.




Fig. 8. Pictures of the Oscilloscope (left) and Synchrowave (right) screen at about 620 seconds into the test, which is marked as point 2 in Fig. 6.

Figure 10 shows pictures of the oscilloscope and the Synchrowave screen at about 870 seconds into the test. At this point, the spoofer has reached its final velocity of 1000 m/s. A phase angle offset of 10° has also been introduced in a matter of minutes. As expected, there is a marked difference in the phasors on the Synchrowave screen. The oscilloscope also shows a time offset of $400 \mu\text{s}$ has been induced in the victim receiver.

Figure 11 shows pictures of the oscilloscope and the Synchrowave screen at about 1370 seconds into the test. At this point, the spoofed signal was heavily attenuated and instantly realigned

with the authentic signals. This was intended to be the end of the test, but when this particular receiver lost lock on the signal it continued to send out a valid time signal to the PMU while fly-wheeling off its internal clock. This caused an alarm to issue on the front panel of the time reference receiver indicating loss of GPS signal lock. The downstream PMU, however, was oblivious to this loss of lock. This state persisted for about half an hour before the clock finally reacquired the authentic signal and instantly realigned its time output, which caused the phasors to realign. Figure 6 does not show the phase angle data for this




Fig. 9. Pictures of the Oscilloscope (left) and Synchrowave (right) screen at about 680 seconds into the test, which is marked as point 3 in Fig. 6.

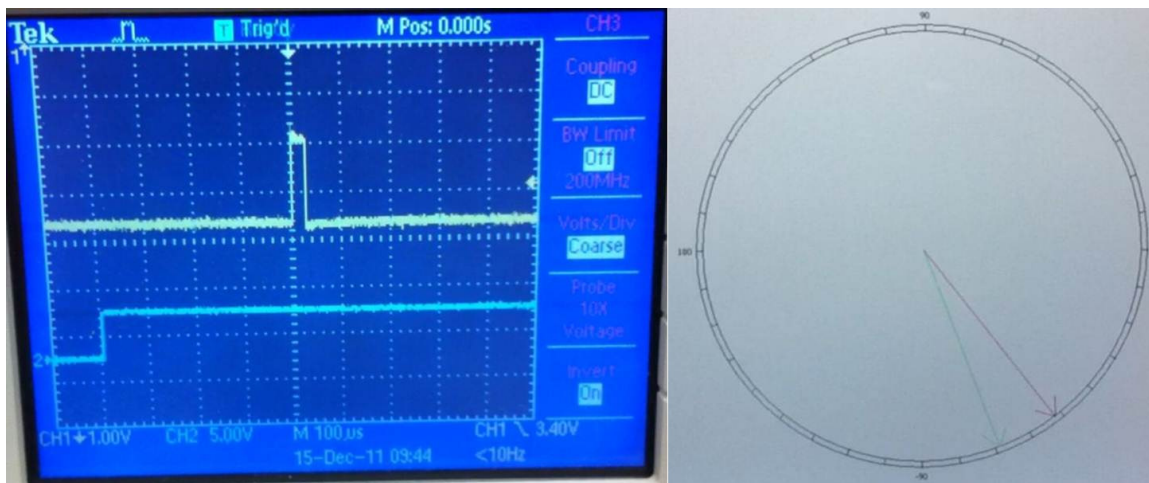


Fig. 10. Pictures of the Oscilloscope (left) and Synchrowave (right) screen at about 870 seconds into the test, which is marked as point 4 in Fig. 6.

entire period, but does show that the phase angle difference exceeds at least 70° before the time reference receiver reacquires the authentic signal.

VI. Implications for Synchrophasor-Based Control

Synchrophasor data provides a clear picture of the state of the power system in real-time. As the size of the power grid grows and stability margins are reduced (to provide more efficient distribution of power), it will become desirable to use synchrophasors for control purposes [9]. PMU manufacturers are currently selling PMUs capable of

implementing automated control schemes that offer response times less than 4 cycles. Such swift response times are seen as necessary to prevent grid instability or damage to equipment.

Control schemes based on synchrophasors rely on phase angle differences between two nodes as an indicator of a fault condition. One example of a currently operational synchrophasor-based control system is the Chicoasen-Angostura transmission link in Mexico [10]. This transmission line links together large hydroelectric generators in Agostura to large loads in Chicoasen through two 400-kV transmission lines and one 115-kV trans-

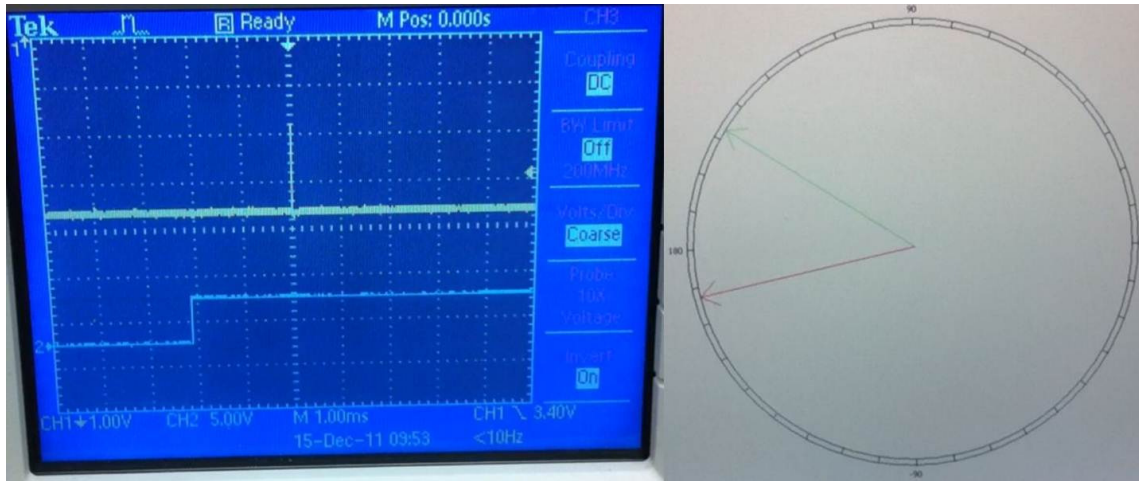


Fig. 11. Pictures of the Oscilloscope (left) and Synchrowave (right) screen at about 1370 seconds into the test, which is marked as point 5 in Fig. 6.

mission line. If a fault occurs in which both of the 400-kV lines are lost, then the hydroelectric generators may experience angular instability. In order to prevent this, a PMU was set up at each end of the transmission lines with a direct communications link between them. It was found that under nominal and single-fault (only one 400-kV line lost) conditions, the phase angle difference between the two locations was less than 7° , whereas a double-fault (both 400-kV lines lost) produced a phase angle difference of 14° . Based on this finding, the PMUs were configured so that if the phase angle difference exceeded 10° the hydroelectric generators would be automatically tripped.

If a spoofer were to attack this system in Mexico or a similar implementation elsewhere, then the spoofer could cause a generator trip. In the test described in the previous section, a 10° offset, the threshold for the Chicoasen-Angostura link, was induced by the spoofer about 250 s after capturing the target receiver, as seen in Figs. 6 and 10. A malefactor could even lead the phase angle off in the opposite direction (say 7°) before cutting both 400-kV transmission lines. Instead of causing a generator to unnecessarily trip, this would prevent PMUs from tripping the generator when required and potentially cause damage to the generator or remaining transmission lines.

Beyond tripping a single generator, there is po-

tential for the effects of the attack to propagate through the grid and cause cascading faults across the grid. One example of this type of cascading failure is the 2003 Northeast Blackout. Although this blackout did not involve PMUs or a spoofing attack, it demonstrates how an appropriately targeted attack against PMUs used for control on the power grid could cause large scale blackouts that originate with a single generator or transmission line trip. On Aug. 14, 2003 at 3:05 p.m., a 345-kV transmission line in Ohio began to sag from increased flow of electric power. When the line sagged too close to a tree, it caused a short-to-ground and tripped offline. This is something that happens fairly frequently on the massive U.S. electrical grid and is usually easily dealt with. However, the tripping of that line in northern Ohio began a cascade of failures that, in a little more than an hour, led to a near total power loss for more than 50 million people in the northeastern U.S. and parts of Canada. The blackout is estimated to have cost approximately 6 billion U.S. dollars for only four days of power loss [1]. This led the Department of Energy and the North American Electric Reliability Corporation (NERC) to fund and push for an improved “smart grid” with synchrophasor technology as a major component.

As previously pointed out, PMUs are high-speed, real-time synchronized measurement devices used to diagnose the health of the electricity grid. With

synchrophasor data, electric utilities can use existing power more efficiently and push more power through the grid while reducing the likelihood of power disruptions like blackouts. Synchrophasor measurements are being looked at to reduce the likelihood of false and inappropriate triggers of transmission system circuit breakers that protectively shut down electrical flow and contribute to cascading blackouts. However, GPS spoofing poses a significant threat to these objectives for PMUs and can make synchrophasor-based control the cause for these events instead of the cure.

VII. Conclusions

Test results presented herein demonstrate that GPS spoofing poses a threat to the integrity of synchrophasor measurements. A spoofer can introduce a time offset in the time reference receiver that provides the timing signal for a PMU without having physical access to the receiver itself. This timing offset produces a corresponding phase offset in the synchrophasor data coming from that PMU. It was demonstrated that a PMU could be made to violate the IEEE C37.118 Standard for synchrophasors in about 11 minutes from the start of a spoofing attack.

As PMU usage continues to grow throughout the world, PMUs will increasingly be used for automatic control purposes instead of just grid monitoring. An example of this is a currently operational system in Mexico which automatically trips a generator if the phase angle difference between PMUs at two particular locations exceeds 10° . The tests discussed in this paper demonstrate that a spoofer could cause control schemes such as the one in Mexico to falsely trip a generator. In the presence of other exacerbating factors, this could lead to a cascade of faults and a large scale blackout.

References

[1] "Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations," Tech. rep., U.S.-Canada Power System Outage Task Force, April 2004.

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

[3] 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.

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

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

[6] Martin, K. E., Hamai, D., Adamiak, M. G., Anderson, S., Begovic, M., Benmouyal, G., Brunello, G., Burger, J., Cai, J. Y., Dickerson, B., Gharpure, V., Kennedy, B., Karlsson, D., Phadke, A. G., Salj, J., Skendzic, V., Sperr, J., Song, Y., Huntley, C., Kasztenny, B., and Price, E., "Exploring the IEEE Standard C37.118-2005 Synchrophasors for Power Systems," *IEEE Transactions on Power Delivery*, Vol. 23, No. 4, Oct. 2008, pp. 1805-1811.

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

[8] 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.

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

[10] 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.