Spoofing GPS Receiver Clock Offset of Phasor Measurement Units

Xichen Jiang, Jiangmeng Zhang, Brian J. Harding, Jonathan J. Makela, and Alejandro D. Domínguez-García

***Abstract*—We demonstrate the feasibility of a spoofing attack on the GPS receiver of a Phasor Measurement Unit (PMU). We formulate the attack as an optimization problem where the objective is to maximize the difference between the PMU's receiver clock offset (with respect to the GPS time measured by the onboard satellite clocks) before and after the attack. Since the PMU uses this clock offset to compute a synchronized time stamp for its measurements, an error in the receiver clock offset introduces a proportional phase error in the voltage or current phase measurements provided by the PMU. For the most general case, the decision variables in the optimization problem are the satellites' ephemerides, pseudoranges, and the receiver's Earth-Centered Earth-Fixed (ECEF) coordinates. The constraints are cast such that the decision variables and the satellite positions computed from the solution of the optimization problem are close to their pre-attack values, so as to avoid possible detection schemes that check for large jumps in the values of these variables. We show that the spoofing attack is feasible for any number of visible satellites. As an illustration of the impact of such spoofing attacks, we present simulation results in which the attack induces errors in a real-time voltage stability monitoring algorithm that relies on the phase information from measurements provided by PMUs.**

***Index Terms*—Phasor measurement unit (PMU), Global Positioning System (GPS), Spoofing.**

I. Introduction

Under the US-DOE Smart Grid vision [1] and its European counterpart [2], electric power systems are undergoing radical transformations in structure and functionality. These transformations are enabled by the integration of new technologies. One such technology that has received considerable attention is the phasor measurement unit (PMU) because of its potential use in real-time monitoring and control (see, e.g., [3] and the references therein). PMUs use the Global Positioning System (GPS) to establish synchronized positive sequence phasor voltage and current measurements [3]. In essence, by acquiring signals transmitted by the GPS satellites and decoding their data, the GPS receiver of a PMU estimates its own position and the offset of its clock with respect to the GPS time measured by the onboard satellite clocks. This essentially enables all the PMUs across a wide geographical area to synchronize their clocks and derive a COORDINATED Universal Time (UTC) time stamp reference.

The authors are with the Department of Electrical and Computer Engineering of the University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA. E-mail: {xjiang4, jzhang67, bhardin2, jmakela, aledan}@ILLINOIS.EDU.

This work was supported in part by the Trustworthy Cyber Infrastructure for the Power Grid (TCIPG) under US Department of Energy Award DE-OE0000097.

The first comprehensive assessment of the vulnerabilities in the civilian GPS infrastructure was published in 2001 by the Volpe National Transportation Systems Center [4]. This report concluded that among the different types of attacks, GPS spoofing is the most malicious and difficult to detect. As defined in [5], spoofing is an intentional interference that misleads a GPS receiver into tracking counterfeit GPS signals. One plausible spoofing method uses a GPS simulator to generate rogue GPS signals matching the genuine signals' phase, code delay, and encoded data [5]. The spoofer gradually increases its transmission power until the GPS receiver locks onto the malicious signals, at which point the victim receiver is fully under the spoofer's control. While major GPS receiver manufacturers have been warned about the lack of effective countermeasures against civilian GPS receiver spoofing, little has been done to address such deficiencies in security [5]; furthermore, most of the civilian GPS receivers on the market today do not have the capability of detecting these spoofing attacks.

Since PMUs use GPS signals to derive a Coordinated Universal Time (UTC) time stamp, they are vulnerable to spoofing. In particular, a spoofing attack can cause the GPS receiver of a PMU to compute an erroneous clock offset, resulting in an erroneous time stamp calculation, which in turn introduces an error in the PMU's phase measurement. Applications in power systems that rely on PMU measurements that may be vulnerable to GPS spoofing include i) health monitoring algorithms (see, e.g., [6], [7]), ii) automatic closed-loop control systems (see, e.g., [3, Ch. 8] and the references therein), and iii) remedial action schemes (see, e.g., [3, Ch. 9] and the references therein). For example, the authors of [8] present a method in which the phase angle measurements from the PMUs are used for wide area stability control of the system; the method provides rapid generator tripping and reactive power compensation switching for transient stability and voltage support. The authors of [9] provide remedial action schemes that determine the severity of a system disturbance by computing the velocity of the phase angles measured by the PMUs and if the velocity is greater than the critical velocity, tripping of generators or transmission lines is initiated.

In this paper, we investigate the feasibility of an attack on the GPS receiver of a PMU. We formulate the attack as an optimization problem where the objective is to maximize the difference between the PMU's receiver clock offset (with respect to the GPS time measured by the onboard satellite clocks) before and after the attack. We perform the optimization for a given instant in time, which is the time when the spoof is to be implemented. The decision variables in this

optimization problem are the satellites' ephemerides,[1] pseudoranges,[2] and the receiver Earth-Centered Earth-Fixed (ECEF) coordinates. Additionally, in order to capture the possibility that the GPS receiver may implement some form of spoofing detection scheme (we elaborate further later in the paper), in the aforementioned optimization problem, constraints are placed on the values that all the decision variables can take.

The problem of spoofing the GPS receiver of a PMU has also been recently addressed in [11]. In this paper, the authors propose an attack method that consists of time shifting the satellites' signals through a delay without making any changes to the data encoded in the signal. This form of spoofing, referred to as a *replay attack*, causes the receiver to overestimate the time of signal transmission, resulting in an incorrect receiver clock offset calculation. In contrast, our spoofing attack method relies on modifying the encoded data without modifying the underlying signal characteristics (although later in the paper, we show that our formulation also encompasses the replay type attack like the one proposed in [11]). As such, we do not expect to be bound by the bandwidth of the receiver tracking loops, but by the rate at which the GPS receiver incorporates new ephemerides and the rate at which the PMU updates its time stamp based on the timing output of the GPS receiver. As we show through numerical examples, an advantage of our proposed data-level spoofing scheme is that, for an arbitrary number of spoofed satellite signals, our scheme also has the additional advantage of being able to introduce a receiver clock offset error without significantly changing the computed receiver position from its pre-attack value, which is important for avoiding detection. On the other hand, with the spoofing scheme in [11], it is not possible to keep the computed receiver position close to its pre-attack value with only a few of the visible satellites' signals compromised because the spoofer can only control the signals' delay but not their data content.

For the specific spoofing attack method presented in this paper, we show that the error introduced in the receiver clock offset can be as high as 2.3 ms, which corresponds to 14% of a cycle for a 60-Hz sinusoidal signal. As an illustrative example, we study the impact of the proposed spoofing method on health-monitoring algorithms that depend on the PMUs' phase measurements. In particular, we show how the introduced phase errors can cause false alarms or misdetections on the estimation of voltage stability margins provided by algorithms that rely on the computation of Thévenin equivalents from measurements provided by PMUs [12]. Another instance of the impact on PMUs from GPS receiver spoofing is illustrated in [13], where the authors showed that alterations to the PMU's receiver clock offset can hamper the performance of algorithms for fault location and oscillation mode monitoring; however, the work in [13] stops short of discussing how to alter the receiver clock offset through a spoofing attack.

[1]The ephemerides are a set of values broadcast by the GPS satellite that allow the receiver to compute the satellite position at a particular time [10].

[2]The pseudorange is the measured distance from the satellite to the receiver and is computed by multiplying the signal propagation velocity (assumed to be the speed of light), by the signal transit time, which is derived from the satellite clock and the receiver unsynchronized clock [10].

The remainder of this paper is organized as follows. In Section II, we describe the algorithm that a GPS receiver uses to compute its position and time offset. In Section III, we formulate the proposed spoofing attack method; while in Section IV, we illustrate its impact on an algorithm for voltage stability monitoring. Section V provides ideas for designing systems to detect such attacks. Concluding remarks are presented in Section VI.

II. Preliminaries

In this section, we explain how a GPS receiver computes its position and clock offset given the satellites' ephemerides and pseudoranges. The relation between the satellite's ephemerides and satellite position is also explained.

A. GPS Receiver Position and Time Synchronization Error

Given the satellites' positions and the times at which the signals containing the data used to compute these positions were broadcast, the GPS receiver can compute its location through a process known as trilateration [14]. Through this process, the receiver computes an estimate of its distance from the satellite by taking the difference between the time of signal transmission and reception and multiplying that by the propagation speed, which is assumed to be the speed of light. If the satellite clock and the receiver clock were perfectly synchronized, this computation yields the true satellite-to-receiver range, and three satellites would be sufficient for the receiver to compute its Earth-Centered Earth-Fixed (ECEF) coordinates. In reality, this is not the case as the receiver clock has an offset, t_u, from the GPS time, t_E, that arise from the internal hardware bias of the receiver clock oscillator.[3]

The receiver clock offset, t_u, is not known a priori by the receiver; thus, the quantity that the the GPS receiver computes by multiplying the estimated time of GPS signal transmission (derived from the unsynchronized receiver and satellite clock) and the propagation speed is not the true range between the satellite and the receiver; instead, this quantity that the receiver computes is referred to as the *pseudorange* [10]. Thus, in order for the GPS receiver to compute its position (in ECEF coordinates), it also needs to estimate the unknown t_u; this means that the receiver needs to track at least four visible satellites. Next, we summarize how the receiver handles the cases of four visible GPS satellites and more than four visible satellites.

1) Four visible satellites: For a given time, let ρ_i and r_i be the i^{th} satellite's pseudorange and true range, respectively, x_i, y_i, and z_i be the i^{th} satellite's ECEF coordinates, x_u, y_u, and z_u be the receiver's ECEF coordinates, $c = 299792458$ m/s, and t_u be the receiver clock offset. Then,

$$\rho_i = r_i - ct_u, \quad i = 1, 2, 3, 4, \tag{1}$$

[3]By assuming that the receiver clock offset t_u is constant across all receiver channels, the GPS time can be expressed as $t_E = t_r - t_u$, where t_r denotes the receiver clock time [14]. The Coordinated Universal Time (UTC), denoted by t_{UTC}, which is used for PMU time synchronization [15], is offset from GPS time, t_E, by an integer number of leap seconds, Δt_{UTC}, i.e., $t_{UTC} = t_E - \Delta t_{UTC}$; as of July 1, 2012, $\Delta t_{UTC} = 16$ s.

where

$$r_i = \sqrt{(x_i - x_u)^2 + (y_i - y_u)^2 + (z_i - z_u)^2}. \quad (2)$$

The satellite coordinates x_i, y_i, and z_i are computed by the receiver through a set of parameters contained in the GPS signal known as the ephemerides (described in detail later). As stated earlier, for each satellite i, the GPS receiver computes the corresponding pseudorage, ρ_i, by multiplying the estimated GPS signal transit time (derived from the unsynchronized receiver and satellite clock) with the speed of propagation, c. In the four visible satellite case and assuming no noise in the measurements and a propagation medium that is vacuum, the receiver location and the clock offset can be obtained by solving (1), as the number of unknowns is equal to the number of equations. This system of nonlinear equations is solved by the GPS receiver through a nonlinear solution method (e.g., Newton-Raphson).

2) More than four visible satellites: When there are more than four satellites visible, as is almost always the case, in (1) we have that $i = 1, 2, \ldots, n$, with $n > 4$, resulting in an overdetermined system. In this scenario, x_u, y_u, z_u, and t_u are obtained by solving a Least Squares Errors (LSE) problem:

$$\underset{x_u, y_u, z_u, t_u}{\text{minimize}} \quad f_0 = \sum_{i=1}^{n} (\rho_i - r_i + ct_u)^2, \quad (3)$$

where $n > 4$ denotes the number of visible satellites. The GPS receiver solves the LSE problem of (3) using a numerical method (e.g., Gauss-Newton).

B. GPS Ephemerides

The ephemerides are a set of parameters that allow the receiver to compute a satellite's position for any time. Up-to-date ephemerides are typically uploaded from the GPS control segment to the satellites once per day and then broadcast to the receiver as part of the navigation data signal. A description of the ephemerides and their role in calculating a satellite's position is presented next. [Interested readers are referred to [16] for a complete description of the ephemerides and how a GPS receiver uses these ephemerides to compute the position of a satellite.]

An accurate characterization of the GPS satellites' orbits is essential for determining the receiver's position. In the absence of external perturbations (classic *two-body* problem), and given the initial time, the trajectory of a satellite can be specified by six constants of integration known as the *Keplerian elements* [14]. In order to describe a satellite's orbit even more accurately under non-ideal conditions, additional forces acting on the satellite must be considered. These forces include the so-called third-body gravitation from the Sun and the Moon, solar radiation pressure, and the Earth's tidal variations, among others. Although the accelerations from the other perturbing forces are small compared to the gravitational acceleration of the Earth, their effects do add up to significant changes over an extended period of time [14].

While it is still possible to completely characterize the satellite's motion under full perturbation with the Keplerian elements, these parameters are no longer constants but depend on time. A reference time known as the *epoch* is established to characterize the dependence with time of the Keplerian elements. At the *epoch*, the six Keplerian elements are such that they describe the statellite's orbit exactly, but as time progresses, the computed position and velocity vectors deviate from the actual position and velocity vectors [10]. In order to account for these deviations, parameters that characterize how the Keplerian elements change over time are added to the satellite's navigation signal. This expanded parameter set which contains the Keplerian elements is known as the satellite's *ephemerides*; at any given time, the GPS receiver uses these ephemerides to compute the position of the satellite.

Next, we provide a compact description of the functional relation between the ephemerides and the satellite position; this functional relation is later used when formulating our spoofing attack method. Let $\delta_i(j)$ denote the j^{th} ephemeride of the i^{th} satellite and define $\overline{\delta}_i = [\delta_i(1),\ \delta_i(2),\ \ldots,\ \delta_i(m)]^T$ to be the vector that contains the ephemerides broadcast by the i^{th} satellite. Then, we can express the ECEF position of the satellite as a function of $\overline{\delta}_i$ such that

$$\begin{aligned} x_i &= f(\overline{\delta}_i, t), \\ y_i &= g(\overline{\delta}_i, t), \\ z_i &= h(\overline{\delta}_i, t), \end{aligned} \quad (4)$$

where the functions $f(\cdot,\cdot)$, $g(\cdot,\cdot)$, and $h(\cdot,\cdot)$ can be extracted from the information provided by Table 20-IV of [16].

III. GPS Receiver Spoofing and Impact on the Phase Information Provided by PMUs

Time synchronization across PMUs is crucial for obtaining accurate phasor angle measurements. In this section, we first discuss how errors in the GPS clock offset arising from a spoofing attack affect PMU time synchronization, and thus the phase information provided by a PMU. Then, we formulate an optimization problem, the solution of which provides a method to spoof a GPS receiver with the intent of introducing an error in the clock offset. In subsequent developments, $\overline{x}$ denotes a vector, the i^{th} entry of which is denoted by $x(i)$. The pre-attack value of a vector $\overline{x}$ is denoted by $\overline{x}^*$; similarly, the pre-attack value of a real-valued variable y is denoted by y^*.

A. Impact of Clock Offset Errors on PMU Phase Information

A GPS simulator can simulate a rogue GPS navigation data signal and cause the GPS receiver of a PMU to latch onto the new signal by gradually overpowering the authentic GPS signal, thus forcing the receiver to compute an incorrect receiver clock offset. In the following developments, we assume that the maximum receiver clock offset from its pre-attack value is not large enough to cause a phase-wrap of 2π in the phase measurement provided by the PMU. Therefore, for demonstrating the feasibility of an attack on PMU time synchronization (and phase measurements), we seek to maximize the absolute difference between the receiver clock offset, t_u, (post-attack) and its pre-attack value t_u^*. For a 60-Hz signal, the PMU's phase measurement error, ε_θ, is related to the receiver clock offset error as follows:

$$\varepsilon_\theta = 60 \times (t_u - t_u^*) \times 360^\circ. \quad (5)$$

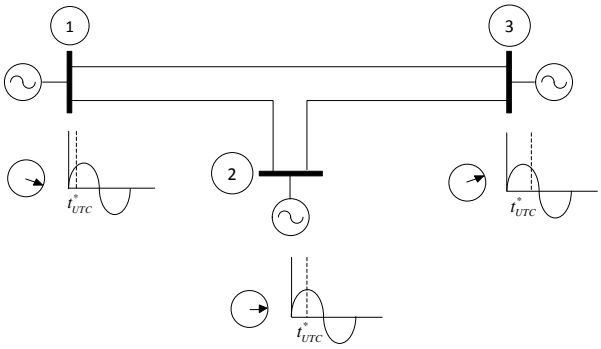

Fig. 1: PMUs and associated phasor measurements.

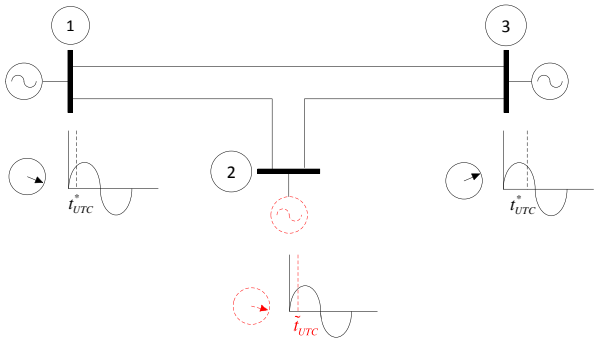

Fig. 2: PMUs and associated post-attack phasor measurements.

Figure 1 conceptually shows a three bus system with PMUs at each bus. A voltage phasor is measured at each bus and time-stamped using the reference time signal t^*_{UTC}. Under normal operation, this time stamp is common to all three buses and provides the synchronization of the PMUs' phasor measurements. Figure 2 shows the result of a GPS spoofing attack on bus 2 of the three bus system. The receiver clock offset, t_u, is shifted from its pre-attack value of t^*_u, causing a proportional error in the estimate of t_{UTC}. Consequently, the erroneous time stamp, $\tilde{t}_{UTC}$, used by the PMU of bus 2 results in an incorrect phase estimate, which causes the measurements provided by this PMU to loose synchronization.

B. Spoofing Attack Formulation

We are interested in determining the maximum phase error that can be introduced in a PMU's phase measurement by spoofing the GPS signal. To this end, the difference between the spoofed clock offset and the pre-attack clock offset is maximized, while considering the possibility that the GPS receiver may implement some form of spoofing detection scheme. Specifically, we assume that these detection schemes preclude certain variables from being perturbed arbitrarily. In the most general case, this imposes upper bounds on the absolute value of the difference between: i) the actual receiver position (as computed by the authentic GPS signals) and the receiver position computed with the spoofed GPS signals; ii) the actual ephemerides and the spoofed ephemerides; iii) the actual satellite positions and the spoofed satellite positions (as computed using the spoofed ephemerides); iv) the actual pseudorange and the spoofed pseudorange (i.e., caused by a replay attack as in [11]). For example, for the absolute value of the difference of the pre- and post-attack receiver position, a possible upper bound would be such that the change in the post-attack receiver position is below the accuracy level of the position provided by the GPS receiver. Additionally, for the absolute value of the difference of the actual satellite positions and the spoofed satellite positions, a possible upper bound would be such that the change in the computed satellite position after spoofing is below the margin of error of the GPS *almanac*[4]. In the optimization problem, the decision variables are the ephemerides, pseudoranges, and the receiver position.

[4]The *almanac* is a reduced precision subset of the ephemerides used to independently approximate the satellite positions to aid in satellite signal acquisition [10].

1) Four visible satellites: We can obtain the expression for t_u by summing the expressions in (1) and solving for t_u, which results in

$$t_u = \frac{-1}{4c}\sum_{i=1}^{4}(\rho_i - r_i(\overline{s}_i, \overline{s}_u)), \tag{6}$$

where $\overline{s}_i = [x_i,\ y_i,\ z_i]^T,\ \forall i$, and $\overline{s}_u = [x_u,\ y_u,\ z_u]^T$. For clarity, we have explicitly written out the dependence of r_i on the satellite position vector, $\overline{s}_i$, and the receiver position vector $\overline{s}_u$. The relations for x_i, y_i, and z_i in $\overline{s}_i$ are as defined in (4), which depend on $\overline{\delta}_i$. Thus, for the four visible satellite case, the maximization of the clock offset error is given by:

$$\begin{aligned}
\underset{\overline{s}_u, \overline{\delta}_i, \rho_i}{\text{maximize}} \quad & (t_u(\overline{\delta}_i, \overline{s}_u, \rho_i) - t^*_u)^2 \\
\text{subject to} \quad & \rho_i = r_i(\overline{\delta}_i, \overline{s}_u) - ct_u(\overline{\delta}_i, \overline{s}_u, \rho_i), \quad i = 1,2,3,4, \\
& |s_u(l) - s^*_u(l)| \le \varepsilon_{s_u}(l), \quad l = 1,2,3, \\
& |\delta_i(j) - \delta^*_i(j)| \le \varepsilon_{\delta_i}(j), \quad j = 1,2,\ldots,m,\ \forall i, \\
& |s_i(k) - s^*_i(k)| \le \varepsilon_{s_i}(k), \quad k = 1,2,3,\ \forall i, \\
& |\rho_i - \rho^*_i| \le \varepsilon_{\rho_i},\ \forall i,
\end{aligned} \tag{7}$$

where $\overline{\delta}_i$ is a vector containing the i^{th} satellite's ephemerides. The differences between the decision variables ($\overline{s}_u$, $\overline{\delta}_i$, and ρ_i) and their pre-attack values (denoted by *) are bounded by $\varepsilon_{s_u}(l)$, $\varepsilon_{\delta_i}(j)$, and ε_{ρ_i}, respectively. We also bound the change in the satellite position, $\overline{s}_i$, by $\varepsilon_{s_i}(k)$. As discussed earlier, these bounds are specified to demonstrate that the spoofing can still succeed even if the receiver checks for large deviations in the values that these variables take with respect to their pre-attack values as a possible countermeasure to detect spoofing (i.e., one possibility to detect large deviations in the satellites' positions is through the GPS *almanac*). If the receiver does not check for abrupt changes as a way to detect data spoofing, then these bounds can be relaxed to positive infinity. In addition, the algebraic relations in (1) must also be satisfied; hence they are included in (7) as equality constraints.

Example 1 (Four-satellite spoofing): In order to illustrate the ideas discussed above, we simulate a spoofing attack on four satellites. We assume that the perturbation of each of the satellites' ephemerides is limited to $\pm 2\%$ of their pre-attack value and the GPS receiver location is restricted to vary at most 15 m from its pre-attack position. We also do not impose any constraints on the changes in the satellite's position, $\overline{s}_i$

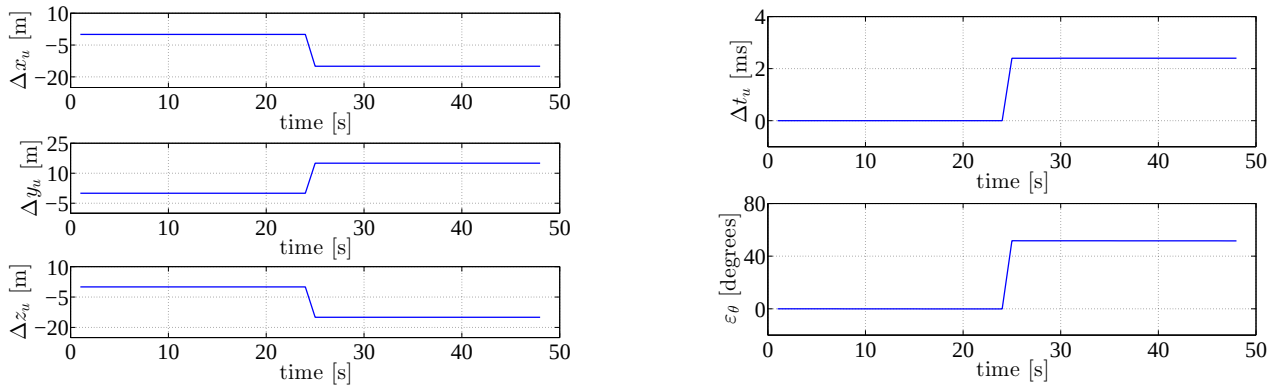

(a) Receiver ECEF coordinates.

(b) Receiver clock offset and phase angle.

Fig. 3: Receiver position, clock offset, and PMU phase error for spoofing four satellites.

(i.e., $\varepsilon_{s_i}(k) = \infty$) and restrict the changes in pseudorange to zero (i.e., $\varepsilon_{\rho_i} = 0$), as to mimic a true data-level spoofing scheme. The optimization problem in (7) is computed for 24 time instances. Solutions for both position and clock offset of the spoofed receiver are plotted along with the corresponding pre-attack solutions in Fig. 3. The attack occurs 24 seconds into the simulation. In Fig. 3(a), it is observed that the jumps in the ECEF coordinates of the receiver due to the spoofed ephemerides are indeed within the 15 m bounds specified by the constraints. Therefore, if the threshold for detecting an attack is greater than 15 m, then such spoofing would not be detected. Figure 3(b) shows the change in the clock offset from the spoofing attack and the resulting phase angle error. ■

2) $n > 4$ four visible satellites: In this case, the system is overdetermined and, in general, an exact solution to (1) no longer exists. Therefore, in the optimization problem of (7), the constraints arising from (1) should be replaced by the LSE condition in (3). Since (3) itself is an optimization problem, it cannot be readily stated as a regular constraint. However, we can exploit this fact and replace the constraints of (1) in the spoofing attack formulation of (7) with the first-order optimality conditions of the LSE problem:

$$\frac{\partial f_0}{\partial x_u} = \frac{\partial f_0}{\partial y_u} = \frac{\partial f_0}{\partial z_u} = \frac{\partial f_0}{\partial t_u} = 0, \tag{8}$$

where

$$\begin{aligned}
\frac{\partial f_0}{\partial x_u} &= 2\sum_{i=1}^{n}\left[\frac{(\rho_i - r_i + ct_u)(x_i - x_u)}{r_i}\right],\\
\frac{\partial f_0}{\partial y_u} &= 2\sum_{i=1}^{n}\left[\frac{(\rho_i - r_i + ct_u)(y_i - y_u)}{r_i}\right],\\
\frac{\partial f_0}{\partial z_u} &= 2\sum_{i=1}^{n}\left[\frac{(\rho_i - r_i + ct_u)(z_i - z_u)}{r_i}\right],\\
\frac{\partial f_0}{\partial t_u} &= 2c\sum_{i=1}^{n}(\rho_i - r_i + ct_u).
\end{aligned} \tag{9}$$

Proceeding as in the four visible satellite case, the variable t_u in the objective function can be solved from any of the expressions in (9); For example, by using $\frac{\partial f_0}{\partial t_u} = 0$, we obtain

$$t_u = \frac{-1}{nc}\sum_{i=1}^{n}(\rho_i - r_i). \tag{10}$$

The problem of maximizing the receiver clock offset when more than four satellites are visible can then be described as

$$\begin{aligned}
\underset{\overline{s}_u,\ \overline{\delta}_i,\ \rho_i}{\text{maximize}} \quad & (t_u(\overline{\delta}_i, \overline{s}_u, \rho_i) - t_u^*)^2\\
\text{subject to} \quad & \frac{\partial f_0}{\partial x_u} = \frac{\partial f_0}{\partial y_u} = \frac{\partial f_0}{\partial z_u} = \frac{\partial f_0}{\partial t_u} = 0,\\
& |s_u(l) - s_u^*(l)| \le \varepsilon_{s_u}(l), \quad l = 1,2,3,\\
& |\delta_i(j) - \delta_i^*(j)| \le \varepsilon_{\delta_i}(j), \quad j = 1,2,\ldots,m,\ \forall i\\
& |s_i(k) - s_i^*(k)| \le \varepsilon_{s_i}(k), \quad k = 1,2,3,\ \forall i,\\
& |\rho_i - \rho_i^*| \le \varepsilon_{\rho_i},\ \forall i.
\end{aligned} \tag{11}$$

Note that in (11), if we relax the bounds $\varepsilon_{s_u}(l)$ and ε_{ρ_i} to positive infinity and restrict $\varepsilon_{\delta_i}(j)$ and $\varepsilon_{s_i}(k)$ to 0 (no data-level spoofing), then we obtain the spoofing attack scheme proposed in [11].

Example 2 (Seven-satellite spoofing): The problem in (11) is solved for $n = 7$ (i.e., seven visible satellites) assuming the same bounds on the receiver position, satellite positions, ephemerides, and pseudoranges as Example 1. The results from the simulation are shown in Fig. 4. The phase angle error resulting from these attacks can be as high as 52°, which corresponds to 14% of a full cycle for a 60-Hz system. Figure 5 shows the receiver clock offset and the resulting PMU phase error for both the four and seven visible satellite spoofing attacks. Comparing the two plots, it can be seen that the maximum phase errors that can be introduced under the same constraints for each satellite are nearly the same. ■

C. Replay Spoofing Attack Formulation

Here, we show how the replay spoofing attack reported in [11] can be cast into the optimization formulation of (11) by restricting the decision variables to only be the pseudoranges, $\rho_i,\ \forall i$. In practice, this corresponds to not changing the ephemeris data contained in the GPS signals. The result is that the signals broadcast by the spoofer are simply delayed

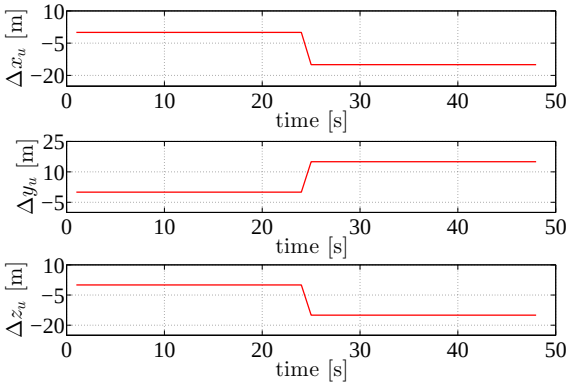

(a) Receiver ECEF coordinates.

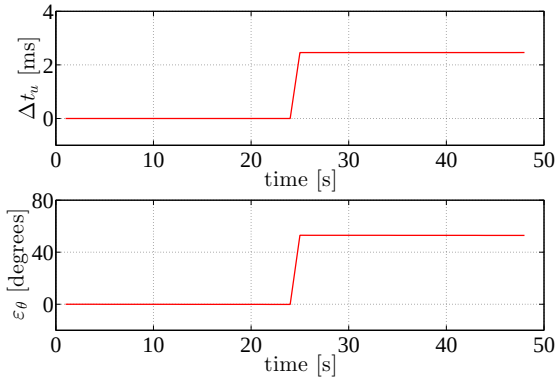

(b) Receiver clock offset and phase angle.

Fig. 4: Receiver position, clock offset, and PMU phase error for spoofing seven satellites.

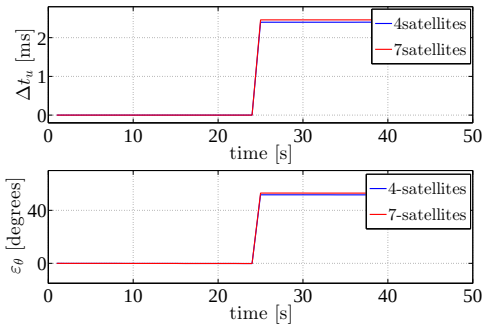

Fig. 5: Clock offset and PMU phase error.

versions of the authentic GPS signals, causing the receiver to erroneously estimate the pseudoranges. [In practice, the success of a replay spoofing attack is also contingent upon hardware limitations such as the receiver's ability to keep lock on the spoofed signal [11].] Then, the problem of maximizing the receiver clock offset can be described by

$$\begin{aligned} \underset{\rho_i}{\text{maximize}} \quad & (t_u(\rho_i) - t_u^*)^2 \\ \text{subject to} \quad & \frac{\partial f_0}{\partial x_u} = \frac{\partial f_0}{\partial y_u} = \frac{\partial f_0}{\partial z_u} = \frac{\partial f_0}{\partial t_u} = 0, \\ & |\rho_i - \rho_i^*| \le \varepsilon_{\rho_i}, \quad \forall i. \end{aligned} \tag{12}$$

Next, we simulate two different replay attacks by using (12) and compare the results with the more general data-level attack method formulated in (11). We show that for an arbitrary number of spoofed satellite signals, our data-level spoofing scheme also has the advantage of being able to introduce a receiver clock offset error without significantly changing the computed receiver position from its pre-attack value. In contrast, a replay attack does not maintain the receiver position close to its pre-attack value when the number of spoofed satellites is small relative to the number of satellites being tracked by the receiver; then, since a PMU is stationary, the receiver may easily detect that its being spoofed by comparing the position it estimates against its a priori known position.

Example 3 (Single-satellite replay attack): Suppose the receiver is receiving signals from seven visible satellites and we spoof one of the seven satellite's GPS signal using the attack method of (12). Without loss of generality, assume that the spoofed satellite is the one that corresponds to index $i = 1$. In (12), we impose $\varepsilon_{\rho_1} = 8000$ m and solve the optimization problem to obtain $t_u = 26.6$ μs. Figure 6(a) shows the receiver's computed position before and after the replay attack, which happens at $t = 24$ s. The receiver's ECEF coordinates all experienced a shift of more than 500 m; thus, if the receiver were checking for such large changes in its position as a spoofing detection scheme, then this attack would not have succeeded. Figure 6(b) shows the resulting receiver

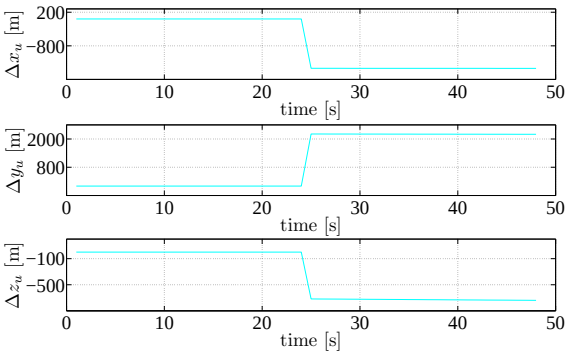

(a) Receiver ECEF coordinates.

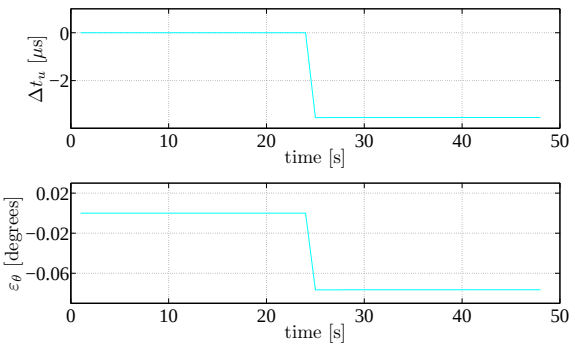

(b) Receiver clock offset and phase angle.

Fig. 6: Receiver position, clock offset, and PMU phase error for replay attack of one satellite.

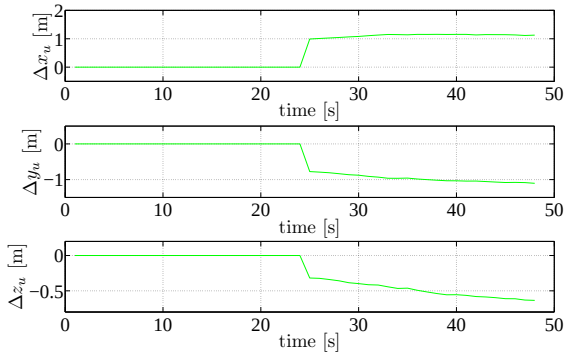

(a) Receiver ECEF coordinates.

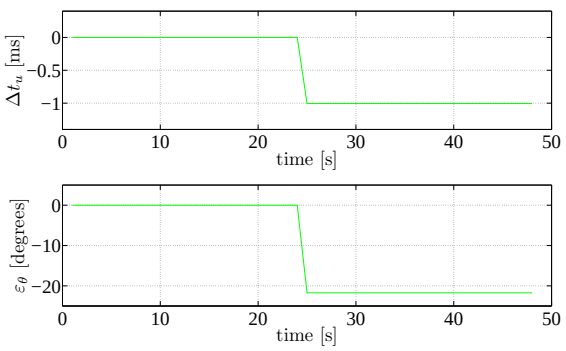

(b) Receiver clock offset and phase angle.

Fig. 7: Receiver position, clock offset, and PMU phase error for replay attack of all seven satellites.

clock offset error and the phase angle error, which is a mere $0.08°$. From the results, we conclude that a replay attack of a single satellite's GPS signal does not introduce significant errors in the receiver clock offset; instead, the errors manifest in the receiver's position. For comparison, we also simulated a data-level spoofing of a single satellite (among seven visible satellites) by solving (11) with the ephemeris of one satellite perturbed. This spoofing attack results in a clock offset error of $t_u = 0.55$ ms, which in turn results in a an error of $12°$ in the PMU phase information. In this case, the receiver position before and after the attack is within 15 m. This lends support that our more general spoofing method can cause significant phase error in the PMU's measurements without forcing the calculated receiver position to change perceptibly from the actual receiver position. ■

In the next example, we show that for a replay spoofing attack to be able to significantly affect the clock offset computation while keeping the GPS receiver position close to its pre-attack value, it is necessary to spoof the signals of several satellites at the same time; this is essentially the spoofing attack method reported in [11].

Example 4 (Seven-satellite replay attack): In this example, we simulate a replay attack on all seven visible GPS satellite signals. In order to obtain a delayed signal by up to 1 ms, we need to allow an error in the pseudorange calculation of up to 300 km with respect to its pre-attack value. Figure 7(a) shows the receiver's computed position, which changes less than 1 m as a result of the attack. If the receiver were checking for large changes in its position (greater than 1 m) as a spoofing detection scheme, then this attack would have evaded spoofing detection. Figure 7(b) shows the receiver clock offset error, which is now 1 ms, and the corresponding phase angle error of more than $20°$. This is large enough to induce misreadings in the stability monitoring algorithms we present later. Hence, in order to introduce significant phase error in the PMUs' measurements without large deviations in the computed receiver position using only a replay attack, all seven satellites must be spoofed together. This introduces additional difficulty as it increases the number of channels required for the GPS spoofer; on the other hand, as illustrated in Example (3), our proposed data-level spoofing attack does not face this restriction. ■

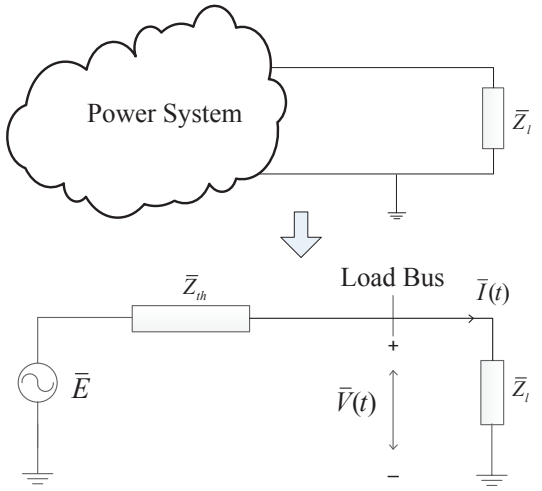

Fig. 8: Thévenin equivalent at load bus.

IV. Impact of GPS Spoofing on a Voltage Stability Monitoring Algorithm

In this section, we consider the impact of GPS spoofing on a voltage-stability monitoring algorithm proposed in [12]; this algorithm relies on the computation of Thévenin equivalents from measurements provided by PMUs. Figure 8 shows a load bus connected to the rest of the power system represented as a Thévenin equivalent; then, in [12], the authors argue that the system is stable if the load impedance magnitude is larger than the Thévenin impedance magnitude. Let $\overline{V}$ denote the voltage phasor at the load bus and $\overline{I}$ denote the current phasor into the load. From two pairs of PMU measurements, and applying KVL, the Thévenin impedance, $\overline{Z}_{th}$, can be computed as

$$\overline{Z}_{th} = \frac{\overline{V}(t_2) - \overline{V}(t_1)}{\overline{I}(t_1) - \overline{I}(t_2)}. \tag{13}$$

Now suppose that a spoofing attack on the PMU's GPS receiver introduces a phase error of ε_θ in the second pair of phasor measurements. The post-attack measurements can be related to the pre-attack measurements as follows:

$$\begin{aligned} \tilde{\overline{V}}(t_2) &= \overline{V}(t_2)e^{j\varepsilon_\theta}, \\ \tilde{\overline{I}}(t_2) &= \overline{I}(t_2)e^{j\varepsilon_\theta}. \end{aligned} \tag{14}$$

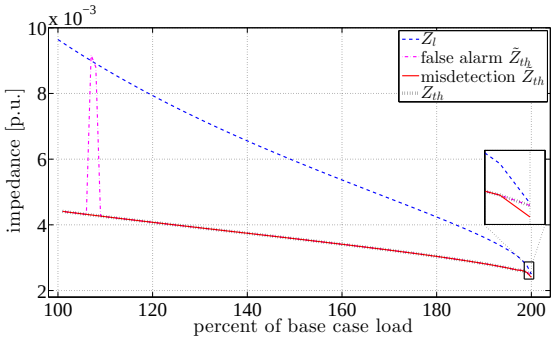

Fig. 9: Pre- and post-attack load and Thévenin impedances.

Then, the post-attack Thévenin impedance, $\tilde{\overline{Z}}_{th}$, is given by

$$\tilde{\overline{Z}}_{th} = \frac{\tilde{\overline{V}}(t_2) - \overline{V}(t_1)}{\overline{I}(t_1) - \tilde{\overline{I}}(t_2)} = \frac{\overline{V}(t_2)e^{j\varepsilon_\theta} - \overline{V}(t_1)}{\overline{I}(t_1) - \overline{I}(t_2)e^{j\varepsilon_\theta}}. \tag{15}$$

We show next how this phase error, ε_θ, can induce false alarms or misdetections in voltage-stability margin estimations.

Consider the standard 3-machine, 9-bus WECC system model (see [17] for a description and parameters). We increase the loads from the base power flow solution in increments of 1% and compute the corresponding Thévenin and load impedance magnitudes at bus 5, denoted as Z_{th} and Z_l; the results are displayed in Fig. 9 with dashed (red) and dotted (blue) plots, respectively. As evidenced from the figure, the magnitudes of the Thévenin and load impedances are approximately equal as the system approaches instability, which is supported by the maximum power transfer theorem. Now suppose that with a loading of 108% of the base case load, we spoof the phasor angles of the PMU's measurements by $-10°$; the resulting Thévenin impedance magnitude, denoted as $\tilde{Z}_{th}$, is then computed to be 0.0092 p.u. This case is shown in Fig. 9 by a dash-dot (magenta) line. Notice that with a shift of $-10°$, the computed Thévenin impedance magnitude is greater than the load impedance magnitude, causing an instance of voltage instability false alarm. When loading is 200% of the base case load, we shift the phasor angles by $2°$, resulting in a calculated Thévenin impedance magnitude of 0.0024 p.u., as shown by the solid (red) line in Fig. 9. We see that the spoofed $\tilde{Z}_{th}$ is much lower than Z_l, giving a false sense of stability when in actuality, Z_{th} is already within the instability region.

Next, we fix the loading and vary the shift in the phase angles of the PMU's measurements in order to see how the computed Thévenin impedance varies with changes in the phase angles. The results are shown in Fig. 10. With a base case loading of 108%, we vary the phase angle and compute the corresponding Thévenin impedances. When the phase angle shift is $0°$, there is no error in the measurements, and the value of the Thévenin impedance magnitude is correct. When the phase angle shift is below $0°$, the Thévenin impedance jumps to value greater than the load impedance, resulting in a voltage instability false alarm. Figure 10(b) shows the same analysis performed with a loading of 200% of the base case load. From the analysis, phase angle spoofing within a range of $(0° - 4°)$ lowers the Thévenin impedance, which does not properly reflect the system's stressed condition.

V. Possible Countermeasures

The effectiveness of the GPS spoofing attacks demonstrated in this paper highlights the need of spoofing detection by PMU receivers. Due to the vast nature of this subject and the rapid developments of recent literature, a thorough treatment of this topic is well beyond the scope of this paper; a good review paper to start with is [18]. Instead, in this section, we provide a broad overview of some of the major techniques currently being developed for possible spoofing detection.

The authors of [19] and [20] recommend the following countermeasures: i) amplitude discrimination, ii) time-of-arrival discrimination, iii) polarization discrimination, iv) angle-of-arrival discrimination, v) cryptographic authentication, and vi) signal strength discrimination. Methods i) and ii) can be implemented in software but only provide a rudimentary defense against spoofing attacks. Methods iii) and iv) require multiple antennas to implement and are ineffective against sophisticated attacks involving multiple rogue GPS transmitters, as discussed in detail in [21]. An extensive review of cryptographic authentication techniques (method v)) can be found in [22]; however, cryptographic methods require significant changes to the current GPS signal coding scheme, which is unlikely to happen in the short term [5]. Recent developments in cryptographic methods, such as navigation message authentication (NMA) and signal authentication sequences (SAS) [23], [24], [25], allow for minimal modifications to the current system. These schemes are robust against signal spoofing but provide no security against unauthorized signal access.

The attack proposed in this paper would be easily detected by methods iv) or v) described above, but iv) requires multiple networked antennas, and v) requires changes to the GPS signal architecture. Several other simple spoofing detection schemes would readily detect the attack proposed in this paper. Consider the case when a receiver is connected to the Internet, It could download the most recent ephemerides from the GPS Control Segment to validate the received navigation data. Since the proposed attack relies on spoofing the ephemerides, a cross-check would reveal tampering. Of course, the spoofing attack would not be revealed until the online ephemeris data were updated, creating a window of opportunity for the spoofer to cause damage in the system.

Instead of checking against published ephemeris data, the receiver could compare the received navigation data with the almanac, a reduced-resolution but multi-satellite version of the ephemerides that is continually broadcast by every satellite along with its own navigation data. By comparing a computed satellite position with the position expected from the almanac, an aggressive spoof could be detected. However, conservative spoofers could stay below any particular threshold by tightening the constraints on the optimization problem.

Most GPS clocks do not use the receiver clock offset measurement directly, but rather use it to guide an independent crystal-controlled oscillator. Monitoring the discrepancy between the oscillator and the computed GPS time could reveal tampering. Another spoofing detection scheme takes advantage of the fact that the genuine satellite signals, while less powerful

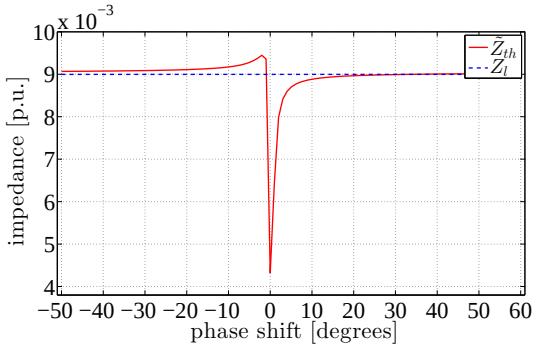

(a) An example of false alarm with 108% of base-case load.

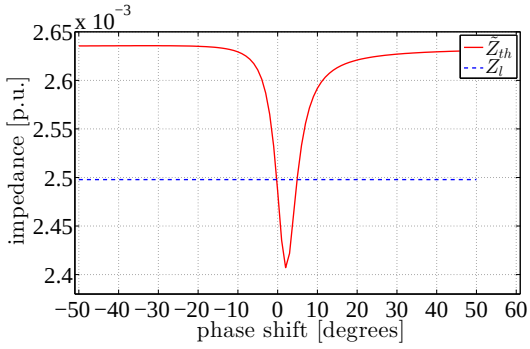

(b) An example of misdetection with 200% of base-case load.

Fig. 10: Thévenin impedances for various phase angle shifts under fixed load values.

than the spoofed signals, are still present. Exact cancellation of the genuine signals would require a complicated spoofer. This technique is known as vestigial signal defense (VSD) and is described in detail in [26]. VSD is software-based and requires no extra hardware. A spoof is detected if additional GPS signals are present in addition to the most powerful ones. The drawback of VSD is that the buried signals are hard to distinguish from multipath interference, but if the GPS receiver is in a static environment (as is the case for PMUs), then multipath effects could be measured and accounted for.

Finally, the proposed spoof would be easily detected in real time if the victim receiver were networked to a trusted GPS receiver at another location. The victim receiver need only validate the navigation data, the current GPS time estimate, or other signal characteristics such as the P(Y) code. The work in [27] shows that spoofing could be revealed by comparing the P(Y) code between the trusted and victim receivers. Since the P(Y) code is an encrypted military code that is transmitted in quadrature with the civilian GPS code, a spoofed signal could not possibly contain a genuine P(Y) code.

VI. Concluding Remarks

PMUs provide synchronized real-time measurements of voltage and current phasors across the power system. They rely on GPS signals to time stamp their measurements and, as such, they are vulnerable to spoofing. In particular, a spoofing attack can cause the GPS receiver of a PMU to compute an erroneous clock offset, which in turn introduces an error in the PMU's phase measurement.

This paper demonstrates the feasibility of an attack on PMU phase measurements through spoofing the ephemerides contained in the GPS signal. An optimization algorithm that maximizes the error in the receiver clock offset while maintaining the computed receiver position close to its pre-attack value is proposed. In the general formulation of the problem, constraints can be added to the decision variables to prevent them from changing significantly as compared to their pre-attack values, since any large, abrupt changes could alarm the receiver of spoofing. When four satellites are visible and the measurements are free of noise, an exact solution to the optimization problem can be found. In the case of more than four satellites, a LSE solution to the optimization problem is formulated with the least squares condition recast into a first order optimality constraint. The feasibility and effectiveness of the proposed spoofing method is demonstrated through multiple simulations of four and seven visible satellite cases.

Subsequent experimental work includes a physical demonstration of this attack on a PMU by building a GPS spoofer. The optimization algorithm will be used to compute the optimal spoofing method for a particular time well in advance of the spoofing attack. The solution of the optimization problem can then be downloaded onto the GPS simulator for execution at a later time. Given the feasibility of such an attack, the effects of erroneous phase measurements must be assessed and countermeasures be developed.

References

[1] "The SmartGrid: An Introduction," 2006. [Online]. Available: http://energy.gov

[2] "European SmartGrids Technology Platform," 2006. [Online]. Available: http://www.smartgrids.eu/documents/vision.pdf

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

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

[5] T. Humphreys, P. Kintner Jr., M. Psiaki, B. Ledvina, and B. O'Hanlon, "Assessing the spoofing threat," *GPS World*, Jan. 2009.

[6] J. Chow, A. Chakrabortty, L. Vanfretti, and M. Arcak, "Estimation of radial power system transfer path dynamic parameters using synchronized phasor data," *IEEE Transactions on Power Systems*, vol. 23, no. 2, pp. 564 –571, May 2008.

[7] H. J. Altuve Ferrer and E. O. Schweitzer III, *Modern Solutions for Protection, Control, and Monitoring of Electric Power Systems*. Pullman, WA: Schweitzer Engineering Laboratories, Inc., 2010.

[8] C. Taylor, D. Erickson, K. Martin, R. Wilson, and V. Venkatasubramanian, "WACS-Wide-Area stability and voltage control system: R d and online demonstration," *Proceedings of the IEEE*, vol. 93, no. 5, pp. 892 –906, May 2005.

[9] S. Stanton, C. Slivinsky, K. Martin, and J. Nordstrom, "Application of phasor measurements and partial energy analysis in stabilizing large disturbances," *IEEE Transactions on Power Systems*, vol. 10, no. 1, pp. 297 –306, Feb 1995.

[10] E. D. Kaplan and C. J. Hegarty, *Understanding GPS: Principles and Applications*. Boston, Massachusetts: Artech House, 2006.

[11] D. Shepard, T. Humphreys, and A. Fansler, "Evaluation of the vulnerability of phasor measurement units to GPS spoofing attacks," in *Proc. of the International Conference on Critical Infrastructure Protection*, Washington, DC, 2012.

[12] K. Vu, M. Begovic, D. Novosel, and M. Saha, "Use of local measurements to estimate voltage-stability margin," *IEEE Transactions on Power Systems*, vol. 14, no. 3, pp. 1029 –1035, Aug. 1999.

[13] Z. Zhang, S. Gong, H. Li, C. Pei, Q. Zeng, and M. Jin, "Time stamp attack on wide area monitoring system in smart grid," in *Computing Research Repository*, Feb. 2011.
[14] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance*. Lincoln, Massachusetts: Ganga-Jamuna, 2011.
[15] "International Earth Rotation and Reference Systems Service," 2012. [Online]. Available: http://maia.usno.navy.mil/ser7/tai-utc.dat
[16] IS-GPS-200D. Interface Specification IS-GPS-200, Revision D, Navstar GPS Space Segment/Navigation User Interfaces, Navstar GPS Joint Program Office, 2004.
[17] P. Sauer and M. Pai, *Power System Dynamics and Stability*. Upper Saddle Rever, NJ: Prentice-Hall, Inc., 1998.
[18] A. Jafarnia-Jahromi, A. Broumandan, J. Nielsen, and G. Lachapelle, "GPS vulnerability to spoofing threats and a review of antispoofing techniques," *International Journal of Navigation and Observation*, vol. 2012, May 2012.
[19] E. Key, "Techniques to Counter GPS Spoofing," Feb. 1995, internal memorandum, MITRE Corporation.
[20] J. Warner and R. Johnston, "GPS spoofing countermeasures," Dec. 2003, Los Alamos Research Paper LAUR-03-6163.
[21] T. Humphreys, B. Ledvina, M. Psiaki, B. O'Hanlon, and P. Kintner, Jr., "Assessing the spoofing threat: Development of a portable GPS civilian spoofer," in *Proc. of ION GNSS*, Sep. 2008.
[22] L. Scott, "Anti-spoofing and authenticated signal architecture for civil navigation systems," in *Proc. of the International Technical Meeting of the Satellite Division of The Institute of Navigation*, 2003.
[23] O. Pozzobon, "Keeping the spoofs out: Signal authentication services for future GNSS," *Inside GNSS*, pp. 48 –55, May 2011.
[24] O. Pozzobon, C. Wullems, and M. Detratti, "Security considerations in the design of tamper resistant gnss receivers," in *Proc. of Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing*, Dec. 2010, pp. 1–5.
[25] O. Pozzobon, L. Canzian, M. Danieletto, and A. Chiara, "Anti-spoofing and open GNSS signal authentication with signal authentication sequences," in *Proc. of Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing*, Dec. 2010, pp. 1–6.
[26] K. Wesson, D. Shepard, J. Bhatti, and T. Humphreys, "An evaluation of the vestigial signal defense for civil GPS anti-spoofing," in *Proc. of ION GNSS*, Portland, Oregon, 2011.
[27] M. Psiaki, B. O'Hanlon, J. Bhatti, D. Shepard, and T. Humphreys, "Civilian GPS spoofing detection based on dual-receiver correlation of military signals," in *Proc. of ION GNSS*, Portland, Oregon, 2011.