Receiver-Autonomous Spoofing Detection: Experimental Results of a Multi-antenna Receiver Defense Against a Portable Civil GPS Spoofer

Paul Y. Montgomery, *Novariant Inc*
Todd E. Humphreys, *University of Texas at Austin*
Brent M. Ledvina, *Virginia Tech*

BIOGRAPHY

Paul Montgomery received a Ph.D in Aeronautics and Astronautics from Stanford University. Paul was a founding member of Novariant, where his accomplishments have included adapting IBLS technology to the Outrider TUAV and X-31 automatic landing systems. In 2006, he was inducted into the Space Technology Hall of Fame for his contributions to the Novariant AutoFarm RTK Autosteer technology. Paul Montgomery is Principal Engineer at Novariant.

Todd E. Humphreys is a research assistant professor in the department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin. He will join the faculty of the University of Texas at Austin as an assistant professor in the Fall of 2009. He received a B.S. and M.S. in Electrical and Computer Engineering from Utah State University and a Ph.D in Aerospace Engineering from Cornell University. His research interests are in estimation and filtering, GNSS technology, GNSS security, and GNSS-based study of the ionosphere and neutral atmosphere.

Brent M. Ledvina is an Assistant Professor in the Bradley Department of Electrical and Computer Engineering at Virginia Tech. He received a B.S. in Electrical and Computer Engineering from the University of Wisconsin at Madison and a Ph.D. in Electrical and Computer Engineering from Cornell University. His research interests are in the areas of software receivers, GNSS applications, estimation and filtering, ionospheric physics, and space weather.

ABSTRACT

In this work we demonstrate the use of a dual antenna receiver that employs a receiver-autonomous angle-of-arrival spoofing countermeasure. This defense is conjectured to be effective against all but the most sophisticated spoofing attempts. The technique is based on observation of L1 carrier differences between multiple antennas referenced to a common oscillator.

We first employ a moderately sophisticated spoofer to "fool" a single-antenna civil receiver. We then deploy the same attack after augmenting the receiver with an additional antenna and with receiver-autonomous spoof-detection software. The work discusses the experimental results together with various issues related to sensitivity, probability of false alarm, impact of carrier multipath, line-bias-calibration, and physical setup and security.

We suggest that this work is important to the community as it provides experimental validation of a low-cost technique for receiver-autonomous spoofing detection. Furthermore, the technique, when combined with physical security of the antenna installation, provides a strong defense against even a sophisticated attack.

The receiver employed is an L1-only civil GPS receiver with multiple antenna capability. The GPS chipset employed is the venerable GP2015/GP2021 that has been freely available for over a decade. As such, this receiver is representative of many civil receivers in use today for a variety of applications. Multiple antennas are enabled either through multiple independent RF front ends and correlators or via antenna multiplexing into a single RF front end and correlator bank.

INTRODUCTION

In 2001, the U.S. Department of Transportation assessed the U.S. Transportation infrastructure's vulnerability to civil GPS disruption [1]. Their report, known as the Volpe report, warned that “as GPS further penetrates into civil infrastructure, it becomes a tempting target that could be exploited by individuals, groups or countries hostile to the U.S.” Among other type of interference, the report considers civil GPS spoofing, an intentional interference whereby a GPS receiver is fooled into tracking counterfeit GPS signals. Spoofing is more

sinister than intentional jamming because the targeted receiver cannot detect a spoofing attack and so can be fooled into generating data that is hazardously misleading.

Previous work into spoofing countermeasures has been carried out, notably in an internal memorandum from the MITRE Corporation in which the author, Edwin L. Key appears to have examined spoofing and spoofing countermeasures in detail [2]. The memorandum recommends the following techniques for countering spoofing:
1.Amplitude discrimination
2.Time-of-arrival discrimination
3.Consistency of navigation inertial measurement unit (IMU) cross check
4.Polarization discrimination
5.Angle-of-arrival discrimination
6.Cryptographic authentication

This paper concerns an implementation of angle-of-arrival discrimination for spoofing detection and employs two or more antennas. Of the proposed techniques, angle-of-arrival discrimination coupled with physical security of the antennas provides significant protection and is relatively easy to implement with inexpensive single frequency receiver technology.

SPOOFING SCENARIOS

Spoofing scenarios can be broadly divided among static and dynamic cases.

Static

A static example is that of a timing receiver that is deployed to synchronize a communications network. GPS time synchronization is also important for the electrical grid and for global trading synchronization. In all such cases, a GPS antenna is situated with clear view of the sky, typically on top of a building or a communications tower. A receiver-generated pulse per second (PPS) is used as the time reference. One can envisage a scenario where the spoofer knows the approximate location of the targeted receiver antenna. Spoofer hardware and a directional antenna could be used to mount an attack at a distance of hundred meters or more. As discussed in [3], the general approach would be:

1. "grow" a replica signal "in the shadow" of the correlation peak for each satellite, replicating the received GPS navigation data
2. increase the power of the spoofing signal to overcome the GPS signal
3. slew the generated signals to be consistent with a desired GPS position/time

Clearly, this technique could be used to fool a timing receiver into generating a PPS that is incorrect.

Dynamic

In fishing waters controlled by the European Commission, large fishing vessels are required to carry a GPS logger [4]. The logger records the voyage of the vessel and provides the data to the licensing agency. Naturally the data can be used to detect passage into waters for which the vessel is not licensed. This is an example of a dynamic scenario where the operator of the vessel is motivated by financial gain to spoof the onboard receiver. In this case of "complicit spoofing" the intent of the operator is to log a fictional voyage that does not disclose illegal fishing activities. In such a situation, the complicit user could potentially disconnect the GPS antenna and attach instead a local GPS signal generator.

There are other dynamic scenarios (largely of Hollywood creation) where a malicious spoofer intends to guide an aircraft into the ground or a mountain side. In our opinion, this would be very difficult to achieve technically and appears to be unlikely in practice. Nonetheless, the mitigation technique presented here could be effective in detecting the presence of a spoofer in these "Hollywood Scenarios."

SPOOFER CATEGORIES

GPS signal generator

Spoofers in this category are GPS signal generators such as can be purchased from several vendors. The RF output is amplified and transmitted, possibly using a directional antenna. In this case, the transmitted signals are not related to the GPS signals being received from satellites in the locality and the navigation data does not replicate the currently active navigation data or timing. Although a receiver could be fooled by this approach, (particularly if the target receiver is first jammed and forced to reacquire), the spoofing signal typically looks like noise to a tracking receiver.

GPS receiver based Spoofer

Spoofers in this category are coupled to a GPS receiver. The GPS receiver tracks satellite signals at a location and decodes the navigation data.

Figure 1: Cornell GRID software-defined GPS receiver

The spoofer then generates a signal that mimics the incident satellite signals in all respects. The spoofer can potentially add a calculated offset to each satellite signal to compensate for a specified geometric offset to the target GPS antenna. The spoofer is also able to vary the signal strength of the constituent signals. This configuration has one transmit antenna, and is moderately sophisticated. A spoofer of this type has been created as discussed in [3]. Although the technical knowledge to create such a spoofer is not widespread the required parts are freely available and may be purchased for a few hundred dollars. Figure 1 shows a photograph of the Cornell "GRID" dual frequency software-defined GPS receiver. As an example of a spoofing platform, the Cornell GRID receiver can simultaneously track 12 C/A channels and generate 8 C/A spoofing channels. Coupled with the simple RF hardware shown in Figure 2, this platform has been used to investigate the challenges involved in mounting a spoofing attack.

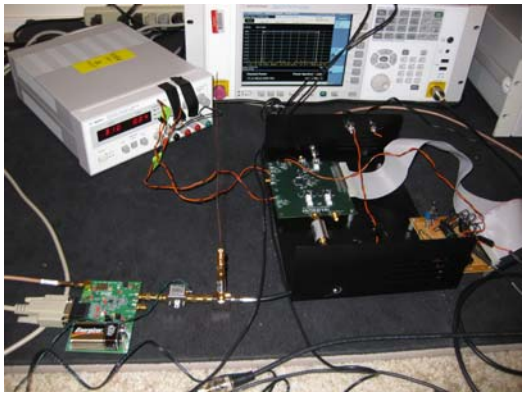

Figure 2: RF transmitter prototype hardware

The hardware described has been used successfully to spoof several different single frequency GPS receivers.

Sophisticated GPS receiver based Spoofer

A sophisticated spoofer design is similar to the GPS receiver based spoofer discussed above, but employs multiple transmit antennas. Furthermore, the spoofer is able to vary the carrier and code phase outputs that are transmitted by each antenna and to control the relative code/carrier phases among these transmit antennas. To create such a spoofer is technically difficult for several practical reasons.

EXPERIMENAL SETUP

Figure 3 shows a photograph of an AutoFarm roof array. This assembly includes a pair of L1 GPS antennas separated by 1.46 meters. Between the antennas is the GPS receiver itself. This assembly was designed and is used for vehicle navigation and automatic control. It is used in the present work because the internal firmware is easily modified and because the assembly produces L1 carrier phase measurements from both antennas referenced to a common internal oscillator. The internal GPS receiver is based on the venerable GP2015/GP2021 chipset and uses Novariant proprietary software. There is nothing special about the hardware platform itself. The platform choice was merely convenient for the authors.

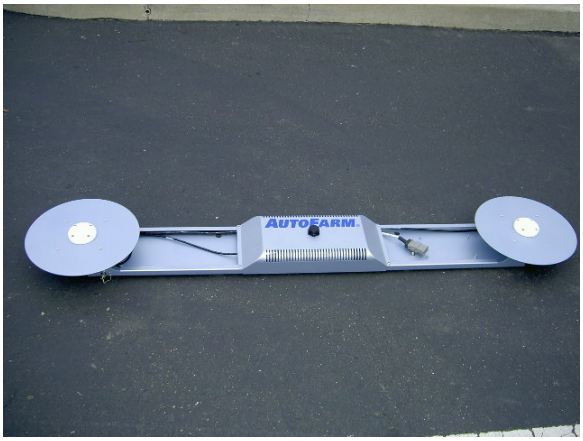

Figure 3: AutoFarm GPS antenna array

In the experimental setup, the assembly is set up on the rooftop in a known and fixed orientation. We chose a level orientation in which the baseline between the antennas is oriented along the (true) North-South axis. Although we are employing a single receiver and common oscillator, the differential phase technique employed is equally applicable to two separate receivers each with a single antenna and oscillator, provided the baseline between the antennas is known.

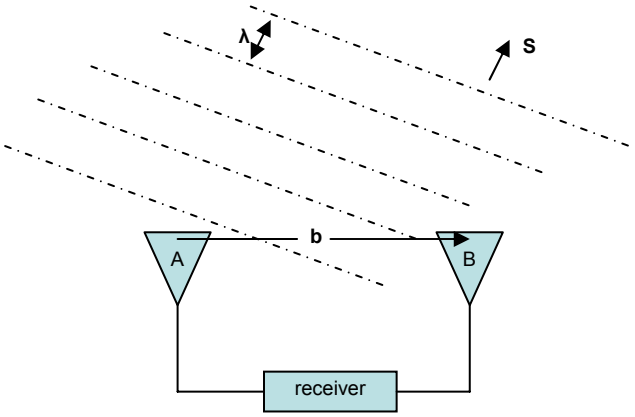

Figure 4: Antenna diversity geometry for a single satellite

Figure 4 shows the geometry for a single satellite:
s is a unit line of sight (LOS) vector to a GPS satellite
b is the baseline vector between the two antenna in units of L1 cycles.
Lines of constant phase emanating from the distant satellite are represented by parallel lines orthogonal to s and separated by the wavelength of the L1 carrier frequency.

For the ith satellite, a scalar equation for the L1 carrier phase difference $d\varphi_i$ between the two antennas is given (in units of L1 cycles) by

$$d\varphi_i = b^T A\hat{s}_i + N_i + B + \gamma_i$$

Equation 1

Where:
b is the baseline vector between the antennas (in the body frame) in units of L1 cycles
A is a direction cosine matrix to rotate vectors in the East-North-Up (ENU) frame to the body frame
$\hat{s}_i$ is the unit line of sight (LOS) vector to satellite$_i$ in the ENU frame
N_i is an arbitrary integer ambiguity for satellite i
B is a constant "line bias" or time varying delta-clock term (depending on implementation)
γ_i is the summation of all carrier phase error terms for satellite$_i$

Equation 1 ignores terms due to the ionosphere and troposphere as these are common mode to sub millimeter level with the assumed meter-level baseline between the antennas. The expression $b^T As_i$ should be recognized as the inner product between vectors s_i and b. For a dynamic case, the LOS vector s_i is known in the ENU frame, but the moving baseline b is known in the body frame. The direction cosine matrix A denotes the attitude of the antenna array and is necessary to compute the inner product. (In other contexts, equation 1 can be used to determine the direction cosine matrix A when it is desired to calculate the attitude of the antenna array.)
In the following, it is assumed that the inner product between s_i and b is known. For a fixed (rooftop) installation the attitude of the antenna array is assumed to be a known by design or pre-survey. For the case of a dynamic (vehicle) installation, the attitude of the antenna array is assumed to be determined by an orientation sensor (such as an inertial attitude sensor) that is not susceptible to GPS spoofing

For a dual-antenna receiver with a common oscillator, B is a constant associated with the difference in the electrical length of the pathways from the antennas to the receiver. For a receiver with separate oscillators (one for each antenna), B is a non-constant bias dominated by the clock offsets between the two oscillators. In either case, B is common for all satellites.

Figure 5 shows a plot of delta phases (dφ) for 4 satellites in track for a period of approximately one hour. The data were generated by the antenna array shown in Figure 3 with the antennas aligned in a North-South orientation. In Figure 5, the integer (N_i) for each satellite has been set to an arbitrary value for convenient plotting. Note that:

- The observed change in dφ over time is due to satellite motion
- The rate of change in dφ is proportional to the baseline length
- Noise on the measurements is due to carrier multipath and carrier phase noise.
- Based on the known attitude expressed in A and the known line bias B, the expected values of the delta phases are plotted 0.1 cycles below the measurement data (for visual clarity). In this case, the attitude was known to approximately 0.1 degrees in pitch and 0.3 degrees in azimuth.

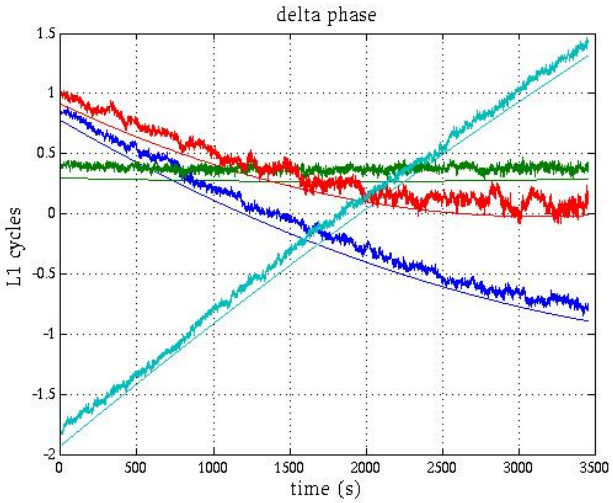

Figure 5: Carrier phase deltas $d\varphi_i$ for 4 GPS satellites over one hour

The plot in Figure 5 illustrates the basic idea for spoofing detection using multiple antennas. If the dφ measurements do not agree with the expected phase profiles within bounds set by the expected noise and attitude uncertainty, then a spoofing signal is identified. As shown in Figure 6, a spoofer transmitting from a single antenna has a very different geometry from that of a receiver tracking GPS satellites distributed across the sky. Consistent with the geometry, the $d\varphi_i$ profiles for a point transmitter must all overlay each other except for contributions due to multipath and phase noise.

The spoofing detection algorithm we implemented is therefore straightforward:

1. The expected delta phases ($d\varphi_i$) are calculated based on known attitude and line bias B.
2. The measured delta phases are compared (modulo 1 cycle) with the expected delta phases.
3. For each satellite, the error between the expected and the measured data is calculated every 500 ms (or other update rate as desired).
4. Based on an error threshold that is a function of satellite elevation, expected worst-case multipath and attitude uncertainty, a limited up-down counter is incremented or decremented.

5. If the up-down counter reaches a specified maximum (based on sample rate and time to alarm) a spoofing detection alarm is triggered.

The algorithm described requires that the attitude of the antenna array be known. This is not a great problem for a static array, however, for the dynamic case it requires integration with an independent attitude reference. An alternate implementation to detect a point transmitter could attempt to detect the case wherein all delta phases overlie each other within an error bound sufficient to accommodate worst-case multipath and carrier noise. Such an approach would have to deal with rare cases where the true satellite geometry happens to cause all delta phases to lie within proximity of each other (modulo one cycle). This situation will occur on occasion and if not handled will lead to a false alarm condition.

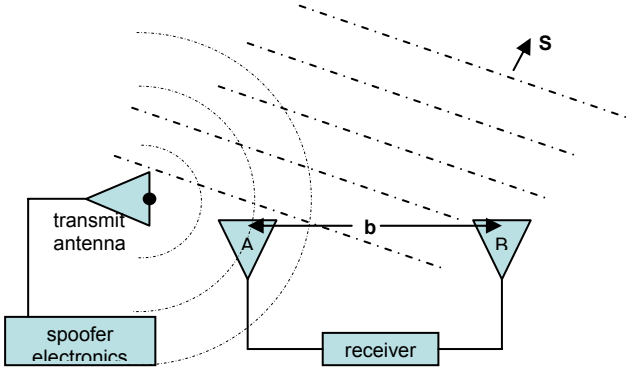

Figure 6: Antenna diversity geometry for a single satellite and point transmitter

For a spoofer to defeat the algorithm as implemented, the spoofing system must emulate the expected carrier phase deltas between the pair of antennas for all satellites in track. As illustrated in Figure 6, it is not possible for even a sophisticated spoofer to emulate this geometry for several satellites with a point transmitter. It is possible that a sophisticated spoofer with two separate points of transmission could defeat the algorithm. However this would also require the spoofer to:

- know the geometry of the GPS antenna array
- locate a matched transmitter antenna very close to each GPS antenna
- deal with other difficult problems associated with multipath, signal leakage and self interference

Such tampering is not possible without physical access to the antenna installation. Of course, the spoofing defense described herein is straightforwardly extended to use 3 or more antennas, making a multi-transmitter spoofing attack even more difficult. For this reason, we believe the use of antenna diversity and physical security leads to a robust defense against even a sophisticated spoofing attack in the case of a static installation.

Figure 7: Pair of CMC AllStar single frequency GPS receivers with independent oscillators

The implementation described above uses a single receiver with a common oscillator. In this case the line-bias (B) is a predetermined constant. If B is not known or non-constant, one can simply remove it by subtracting one measurement from all the others. (typically one subtracts the measurement from a high elevation satellite). With this modification and with a commensurate increase in the error threshold to accommodate increased phase noise and multipath, the algorithm is then easily implemented with separate receivers and antennas. In the multi-receiver case, it is only required that vector between the two antennas is constant and known. Figure 7 shows a pair of low-cost, single frequency CMC AllStar receivers. A pair of similar receivers can be used to provide a timing reference (~50 ns) that is robust to even a sophisticated single-transmitter spoofing attack. In the case where an existing receiver is used to provide GPS timing, it is also possible to upgrade the installation by adding a second receiver and antenna. The algorithm described can then be used to achieve protection against spoofing attack.

INDOOR EXPERIMENT

After the Receiver Autonomous Spoofing Detection (RASD) software was developed, the antenna array depicted in Figure 3 was mounted on the roof and the software enabled. The algorithm was tested for several days in an "unspoofed" setting to validate that spurious (false) detections were not flagged. A detection threshold of +/-0.1 cycles was used for zenith satellites, with a linear increase to +/-0.25 cycles for satellites at the horizon. These thresholds were found empirically.

The intended experiment was to employ the hardware depicted in Figures 1 and 2 to implement a spoofing attack against the GPS receiver and antenna array

depicted in Figure 3 in a situation where the array was mounted to the roof with a clear view of the sky as follows:

1. Create a spoofing attack against the test receiver without the RASD firmware and validate (using an oscilloscope) that the spoofer was able to drive the generated PPS away from truth.
2. Enable the RASD firmware and execute the same attack. This time validate that the receiver was able to detect the attack and raise an alarm.

This plan was complicated by the necessity to perform the experiment outside and with a clear view of the sky. The issue: transmitting outdoors in the L1 frequency band is illegal! Moreover, an early version of the spoofing system in Figures 1 and 2 suffered from a software bug (since resolved) that prevented reliable transmission of the spoofing signal.

In view of these problems, it was decided to simplify the experiment and move indoors. The revised experiment used a "real" GPS signal from a rooftop mounted GPS antenna. The L1 band was amplified and re-transmitted indoors from a point source as illustrated in Figure 8. (This setup is identical to that occasionally used at trade shows to allow GPS receiver vendors to demonstrate active GPS receivers indoors without cable runs to the roof)

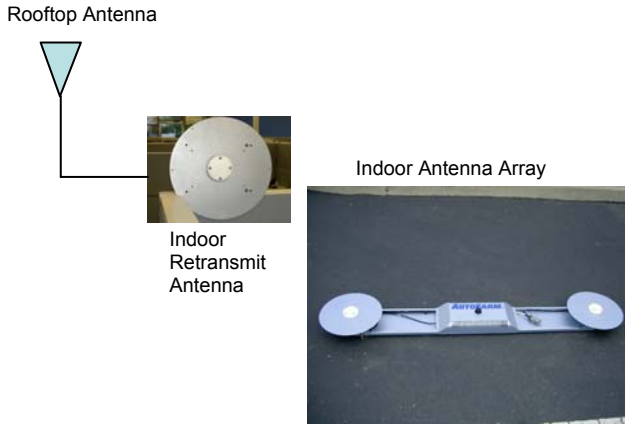

Figure 8: Indoor Experimental Setup

The antenna array was likewise moved indoors and mounted in proximity to the re-transmit antenna. In this environment the receiver was able to track the re-transmitted signal without cycle slips when the range to the transmit antenna was less than approximately 6 meters. The indoor receiver was also able to decode all navigation data and calculate a position fix (the position of the rooftop antenna). Having calculated the position fix and the LOS vectors, the receiver immediately detected and flagged a spoofer signal. The delta phase profiles of seven satellite signals that were tracked during the indoor experiment are plotted in Figure 9. Consistent with the geometry, one may observe that all delta phases from the point source lie on top of each other. Although this experiment did not detect an actual spoofer, the situation is sufficiently representative that the approach and software implementation were considered successful.

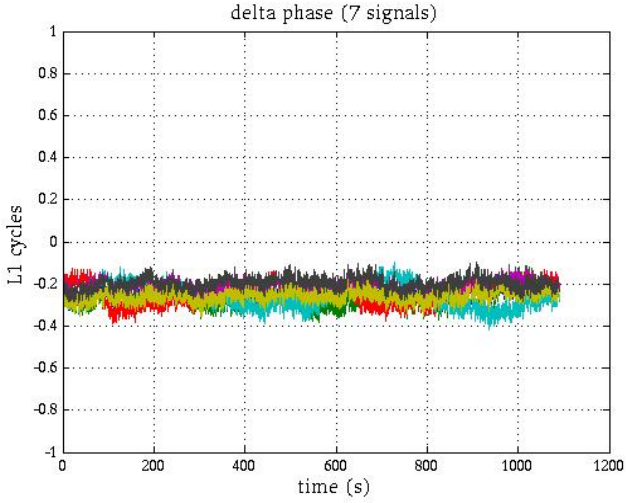

Figure 9: Experimental delta phase results with point transmitter

DYNAMIC PLATFORM

While attempts to spoof a static target to affect timing or position integrity may be a more obvious threat, there are also cases where attempts may be made to spoof a non-static target (recall the foregoing example of the fishing vessel skipper who desires to spoof an onboard GPS-based monitoring unit to fish undetected in forbidden waters).

Figure 10: Crossbow AHRS440 MEMS based Attitude Sensor

Attitude		
Range: Roll, Pitch (°)	± 180, ± 90	
Accuracy[2] (° rms)	< 0.5	With external GPS aiding
Accuracy[2] (° rms)	< 1.5	Without external GPS aiding

In this case the multi-antenna spoofing defense is aided by an independent attitude sensor. A MEMS-based attitude sensor may be sufficient for this purpose if the antenna baselines are kept small. Consider an antenna baseline of 0.36 meters and the Crossbow AHRS440 MEMS based attitude sensor shown in Figure 10. Accepting a 1σ attitude accuracy of 1 degree about each axis, a 3σ attitude error will result in a worst case phase error of approximately 0.1 cycles. Applying alarm bounds widened sufficiently to accommodate this attitude uncertainty in addition to worst-case carrier multipath and phase noise, one concludes that a low cost MEMS attitude sensor could be used in cooperation with multiple antennas to provide spoofing detection for many dynamic situations.

CONCLUSIONS

The use of antenna diversity using either multiple separate receivers or a multi antenna single-oscillator receiver can be used to defend against intentional GPS spoofing. This is because the use of multiple antennas greatly increases the technical difficulty required to mount a successful spoofing attack. In general, an additional spoofer transmitter is required for each additional GPS antenna. Furthermore, a spoofer would have to locate each transmit antenna in close physical proximity to the appropriate GPS antenna in the array. If the GPS antennas of a static installation are further protected by physical security it is possible to create a robust defense against even a sophisticated spoofing attack.

One can broadly divide GPS applications into static and dynamic cases. In static cases it is feasible to construct and survey an antenna array and provide physical security to the antenna array. One time survey of the fixed array is sufficient to enable receiver autonomous spoofing detection (RASD).

In dynamic cases, it is possible to determine the attitude of the antenna array with sufficient accuracy from a non-GPS-based sensor. Inexpensive MEMS-based attitude sensors could provide sufficient attitude accuracy to aid detection of intentional spoofing. For the case of a complicit user, the presence of multiple antennas makes it difficult to intentionally defeat the system by direct injection of an artificial GPS signal. The cost of the required attitude sensor and the probability of false or missed detection may be adjusted by the tradeoff between baseline length and attitude sensor accuracy. The technology to enable multi-antenna spoofing detection is freely and immediately available using any of the numerous GPS receivers that produce L1 carrier phase observables.

ACKNOWLEDGMENTS

The authors would like to thank Novariant for the use of the AutoFarm roof array used for the experiment. Special thanks to Dennis Connor for supporting this work.

REFERENCES

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

[2] Key, E. L., "Techniques to Counter GPS Spofing," Internal memorandum, MITRE Corporation, Feb. 1995.

[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 ION GNSS 2008*, Institute of Navigation, Savanna, GA, 2008.

[4] Private communication with European Commission Joint Regulators' Group, October 2008.