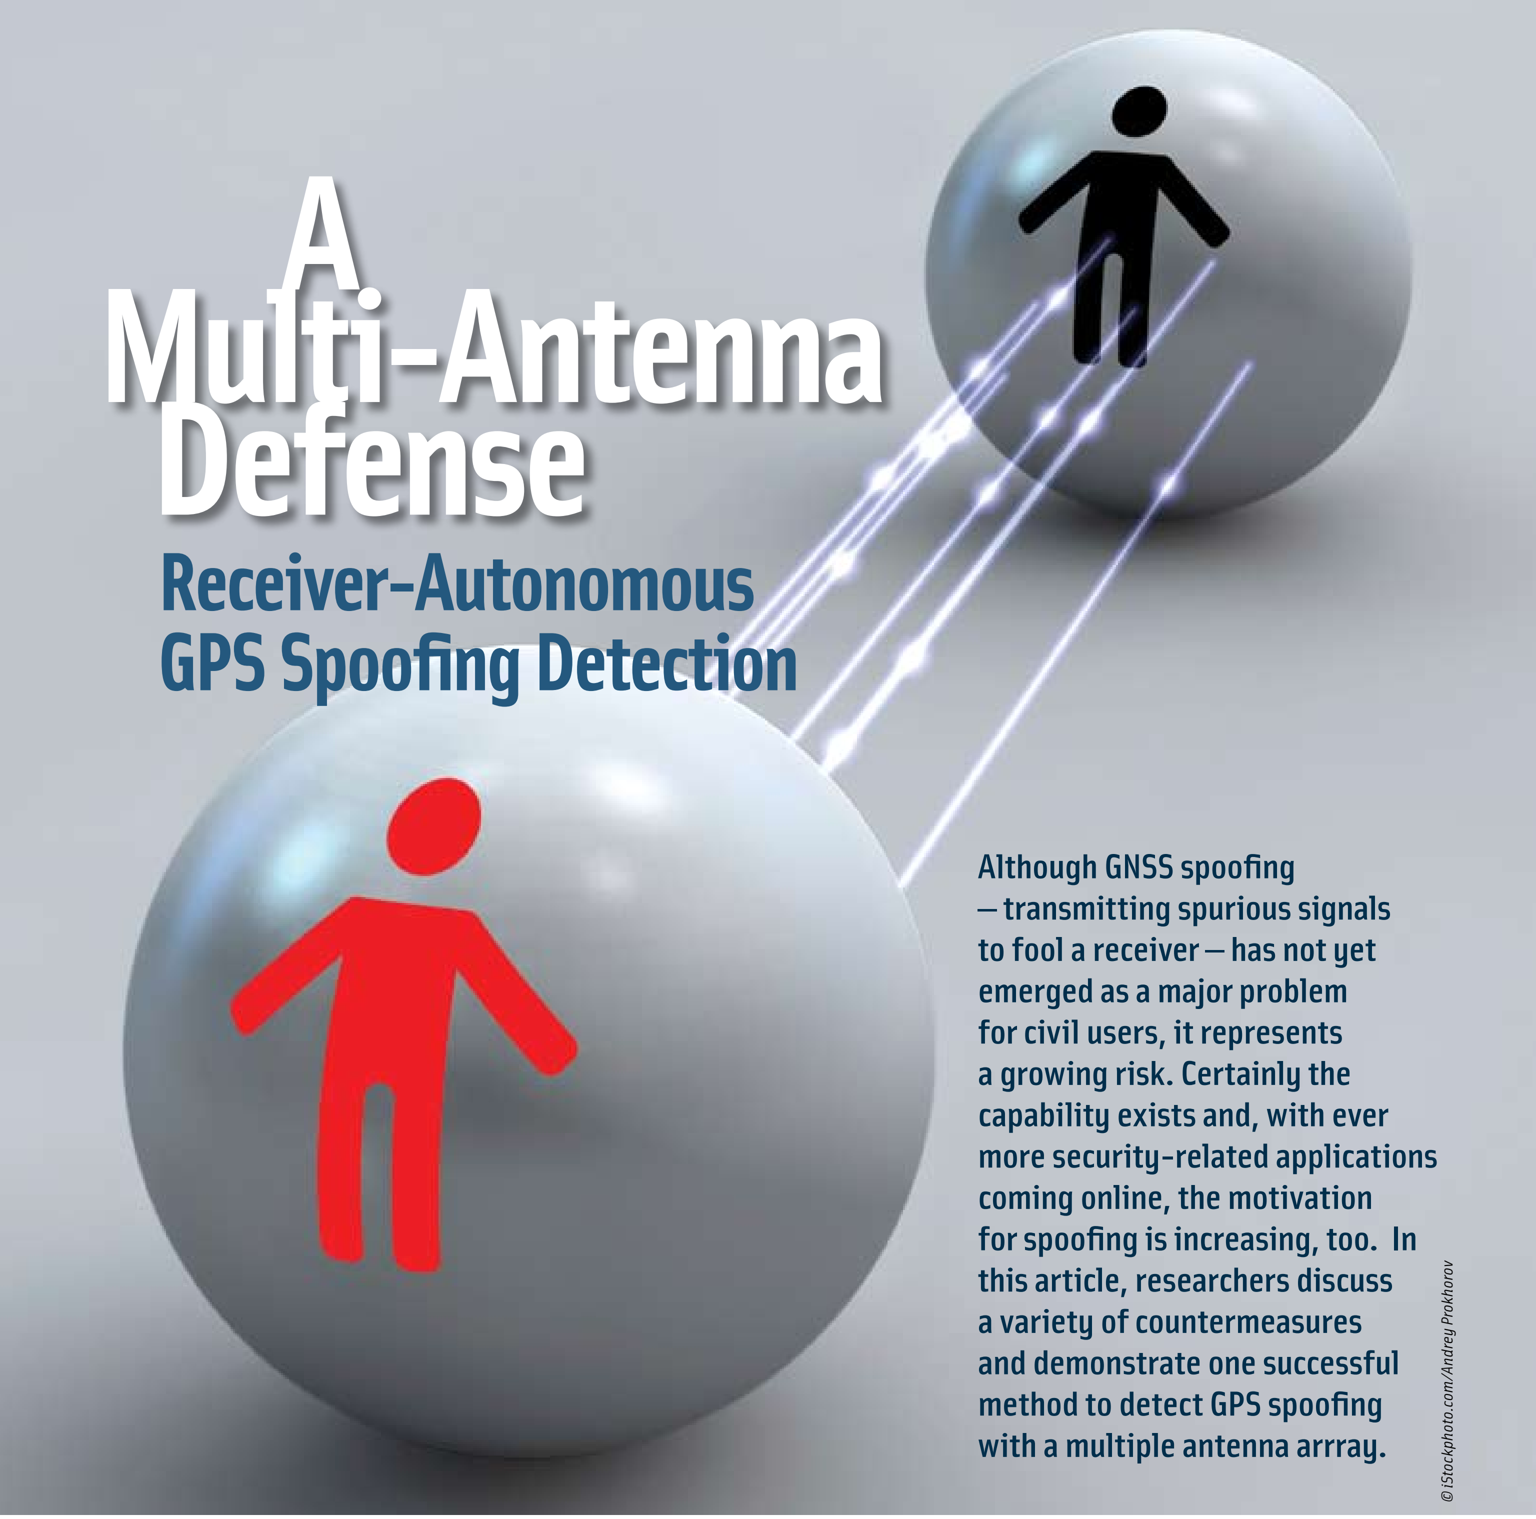

A Multi-Antenna Defense

Receiver-Autonomous GPS Spoofing Detection

Although GNSS spoofing – transmitting spurious signals to fool a receiver – has not yet emerged as a major problem for civil users, it represents a growing risk. Certainly the capability exists and, with ever more security-related applications coming online, the motivation for spoofing is increasing, too. In this article, researchers discuss a variety of countermeasures and demonstrate one successful method to detect GPS spoofing with a multiple antenna arrray.

© iStockphoto.com/Andrey Prokhorov

PAUL Y. MONTGOMERY
NOVARIANT INC
TODD E. HUMPHREYS
UNIVERSITY OF TEXAS AT AUSTIN
BRENT M. LEDVINA
VIRGINIA TECH

The issue of intentional or inadvertent interference to GNSS signals is a matter of growing concern throughout the world.

In a study released the day before the terrorist attacks on the Pentagon and the New York World Trade Center in September 2001, the U.S. Department of Transportation assessed the national transportation infrastructure's vulnerability to civil GPS disruption.

The agency's investigation and subsequent recommendations, 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."

A few years later, in a 2004 National Security Presidential Directive on space-based positioning, navigation, and timing (PNT), former U.S. President Bush gave the Department of Homeland Security (DHS) responsibility for leading development of a plan to address concerns about interference to GPS.

DHS issued a preliminary interference detection and mitigation (IDM) plan last year.

To date, most actual incidents involving GPS interference — whether intentional or unintentional — have involved in-band or out-of-band harmonic RF transmissions that masked the weak GPS spread spectrum signals. A good deal of anxiety has been expressed in recent years about inexpensive GPS jammers that, at power levels as low as one watt, could cause wide areas of disruption to GPS service.

Among other types of intentional interference, the Volpe report and the IDM plan mention civil GPS spoofing, a technique by which a GPS receiver is fooled into tracking counterfeit GPS signals. Spoofing is more sinister than intentional jamming because it is surreptitious: the targeted receiver cannot detect the attack and, consequently, can be fooled into generating erroneous data that may even be hazardously misleading.

Previous work into civil spoofing countermeasures begins with an important internal memorandum from the MITRE Corporation in which the author, Edwin L. Key, appears to have examined spoofing and spoofing countermeasures in detail. (For details, see the "Additional Resources" section near the end of this article.) The memorandum recommends the following techniques for spoofing detection:

1. amplitude discrimination
2. time-of-arrival discrimination
3. consistency of navigation and inertial measurement unit (IMU) cross check
4. polarization discrimination
5. angle-of-arrival discrimination
6. cryptographic authentication

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.

In this article we demonstrate the use of a dual-antenna receiver that employs a receiver-autonomous angle-of-arrival spoofing countermeasure — essentially an implementation of Key's fifth technique. It is based on observation of L1 carrier differences between multiple antennas referenced to a common oscillator. We believe that this defense could be effective against all but the most sophisticated spoofing attempts.

In one example of "complicit spoofing," the intent of the operator would be to log a fictional voyage that does not disclose illegal fishing activities.

Spoofing Scenarios

Spoofing scenarios can be broadly divided into static (fixed target receiver) and dynamic (moving target receiver) cases.

Static Scenario. The target receiver of a static spoofing scenario could be, for example, a timing receiver deployed to synchronize the electrical power grid, global trading, or a communications network.

In all such timing applications, the GPS antenna is situated with a 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 for synchronization.

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 the paper by T. E. Humphreys et alia cited in Additional Resources, the general approach would be as follows:

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 signals to overcome the GPS signals
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 Scenario. Since January 2005, in fishing waters controlled by the European Union (EU), Commission Regulation No. 2244/2003 has required that operators of fishing vessels more than 15 meters in length carry a satellite-based vessel monitoring system .

The VMS (typically employing GPS today) records the voyage of the vessel and automatically provides the data to the fisheries monitoring center of the EU member state where the vessel is registered as well as the member state in whose waters the vessel is fishing.

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. (Indeed, Article 6, Section 2 of regulation 2244/2003 specifies that the VMS data are not to be altered in any way and prohibits anyone from attempting to "destroy, damage, render inoperative or otherwise interfere with the satellite tracking device.")

In this case of "complicit spoofing" the intent of the operator would be to log a fictional voyage that does not disclose illegal fishing activities. In such a situation, the complicit user could disconnect the GPS antenna and attach instead a local GPS signal generator.

This is exactly the scenario envisioned in a November 25, 2008, presentation at a conference on MENTORE (iMplemENtation of GNSS tracking & tracing Technologies fOR Eu regulated domains), an R&D project funded by the European Commission's 6th Framework Program and coordinated by the Institute for the Protection and Security of the Citizen (IPSC) of the EC's Joint Research Centre (JRC).

In the presentation, JRC researcher Gianmarco Baldini referred to the existence of a "simple GPS fraud kit" available for €2,000 (about $2,535) that could feed spoofed signals into the VMS's RS-232 port after disabling the antenna and opening the VMS box. Baldini also

Cornell University GRID software-defined GPS receiver (left) RF transmitter prototype hardware (right)

mentioned "sophisticated GPS signal simulators" available for about €100,000 ($126,700) that could be connected to or radiated towards the VMS unit.

Other dynamic scenarios (largely of cinematic imagination) might involve a malicious spoofer who seeks to guide an aircraft into the ground or a mountainside. 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 such "Hollywood Scenarios."

Spoofer Categories

The JRC MENTORE paper cited earlier mentioned a couple of categories of GPS spoofers. In this article we identify three main types of spoofers.

GPS signal generator. Spoofers in this category are GPS signal generators readily available from several vendors. For use as a spoofer, the signal generator's RF output is amplified and transmitted, possibly using a directional antenna.

In this case, the transmitted signals are not phase- and frequency-matched to the GPS signals being received from satellites in the locality, and the navigation data do not replicate the currently active navigation data. Although a receiver could be fooled by this approach, particularly if the target receiver is first jammed and forced to reacquire, the spoofing signal generated in this fashion typically looks like noise — rather than a usable signal — to a receiver tracking it.

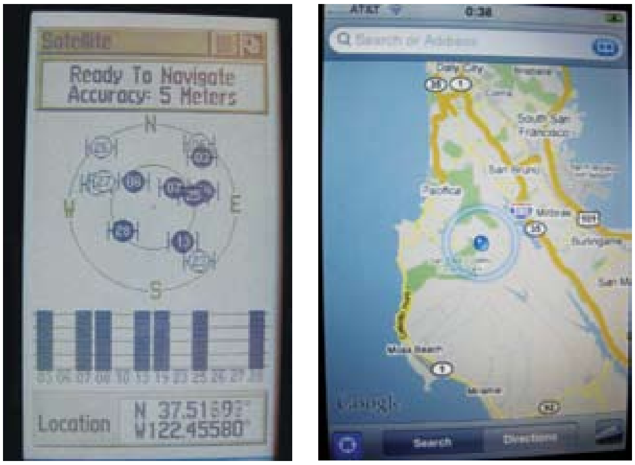

Screen shots of spoofed commercial handheld GPS receivers

GPS Receiver 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.

The spoofer then generates a signal that mimics the incident satellite signals in all respects. Conceivably, a spoofer could 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 so that they appear at the target antenna to have the same relative strengths of the authentic signals.

This spoofing configuration has one transmit antenna and is moderately sophisticated. This kind of spoofer has been created, as discussed in the Humphreys paper cited earlier. 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.

An accompanying photograph shows the Cornell University "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 the second photo, this platform has been used to investigate the challenges involved in mounting a spoofing attack. The hardware has been successful, under controlled laboratory conditions, in spoofing several different single-frequency GPS receivers, as illustrated by the screen shots in the accompanying photos.

Sophisticated GPS Receiver-Based Spoofer. This kind of design is similar to the equipment described in the previous category but employs multiple transmit antennas. Furthermore, the spoofer is able to vary the carrier phase outputs that are transmitted by each antenna to control the relative carrier phases among these transmit antennas. Creating such a spoofer is possible but technically difficult.

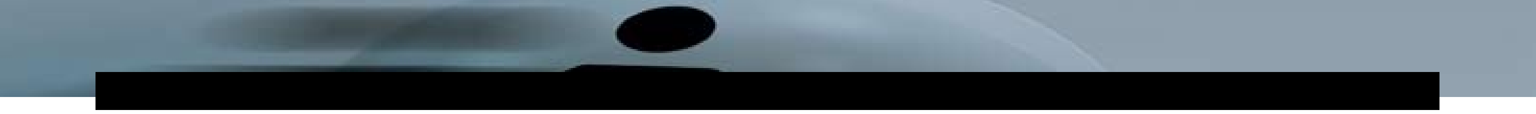

Setting Up the Experiment

To help investigate our spoofing detection technique, we used a dual-antenna array mounted on a rooftop as shown in the accompanying photograph. This assembly includes a pair of L1 GPS antennas separated by 1.46 meters. Between the antennas is the GPS receiver itself.

The receiver/antenna assembly was designed for vehicle navigation and automatic machine control. We used it in this project because the internal firmware is easily modified and the assembly produces L1 carrier phase measurements from both antennas referenced to a common internal oscillator.

GPS antenna array

The internal GPS receiver is based on a 12-channel, L1 C/A-code chipset and uses proprietary software. There is nothing special about the hardware platform itself; the platform choice was merely convenient for the authors.

In the experiment, we installed the assembly on the rooftop in a known and fixed orientation. We chose a level configuration in which the baseline between the antennas is oriented along the (true) north-south axis. Although we used 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.

Methodology for Detecting a Spoofing Attack

Exploiting the equipment configuration that we have described, we developed a technique for detecting spoofing signals based on their deviation from the characteristics of signals received from actual GPS satellite transmissions.

Figure 1 shows the geometry for a single satellite in which

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 *i*th 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 Equation 1

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

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 the submillimeter 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 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 needed 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 discussion, 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 2 shows a plot of delta phases ($d\varphi$) for four satellites tracked for approximately one hour. Note that the various satellites' delta phases vary distinctly from one another. These data were generated by the antenna array described earlier with the antennas aligned in a north-south orientation. In Figure 2, the integer (N_i) for each satellite has been set to an arbitrary value for convenient plotting.

Note that:

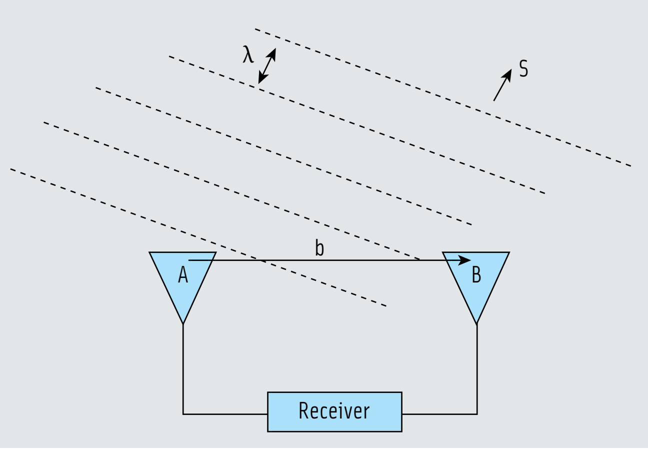

FIGURE 1 Antenna diversity geometry for a single satellite

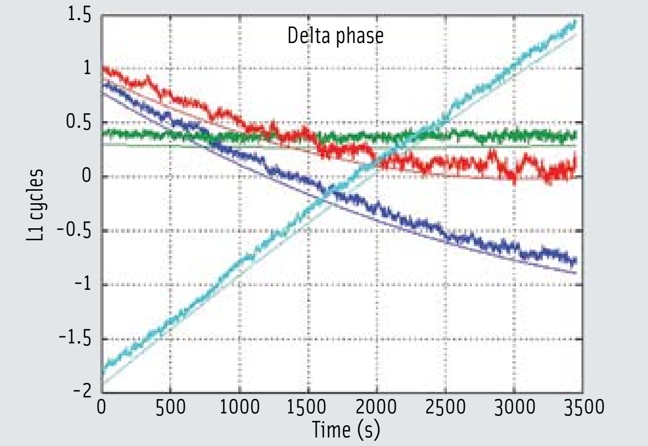

FIGURE 2 Carrier phase deltas di for four GPS satellites over one hour

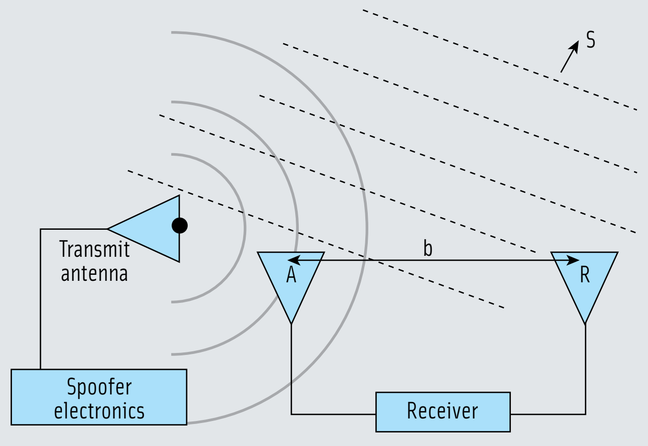

FIGURE 3 Antenna diversity geometry for a single satellite and point transmitter

- 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 shown in Figure 2 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.

Identifying a Spoofed Signal

The plot in Figure 2 illustrates the basic idea for spoofing detection using multiple antennas. If the $d\varphi_i$ 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 3**, 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 would 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 milliseconds (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), it triggers a spoofing-detection alarm.

The algorithm described here requires that the attitude of the antenna array be known. This is not a great problem for a static array; however, the dynamic case requires integration with an independent attitude reference.

An alternative attitude-independent implementation of the multiple-antenna spoofing detection technique could be designed to raise an alarm when the delta phases are suspiciously close to one another; that is, when all the delta phases overlay each other within a bound that is just wide enough to accommodate worst-case multipath and carrier noise.

Such an approach would have to deal with rare cases where, for a brief moment, the true satellite geometry happens to cause all delta phases to overlay each other (modulo one cycle). This situation will occur on occasion and, if not handled correctly, will lead to a false alarm.

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 3, even a sophisticated spoofer cannot emulate this geometry for several satellites if the spoofer is limited to one transmitting antenna.

A sophisticated spoofer with two separate points of transmission might be able to 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, and deal with other difficult problems associated with multipath, signal leakage, and self-interference.

Such an attack would not be possible without physical access to the antenna installation. Of course, we could extend the spoofing defense described here in a straightforward way by using three 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 sophisticated spoofing attacks.

The implementation that we have just described uses a single receiver with a common oscillator. In this case, the line-bias (*B*) is a predetermined constant.

If *B* is unknown or non-constant, one can simply remove it by subtracting one satellite measurement from all the others. (Typically, one subtracts the measurement from a high elevation satellite.)

With this modification and a commensurate increase in the error threshold to accommodate increased phase noise and multipath, the algorithm is easily implemented using separate receivers and antennas. The multi-receiver approach only requires that the vector between the two antennas is constant and known.

Indoor Experiment

After developing the receiver autonomous spoofing detection (RASD) software, we mounted the antenna array depicted earlier on the roof and enabled the software. The algorithm was tested for several days in an "unspoofed" setting to validate that spurious (false) detections were not flagged.

We used a detection threshold of ±0.1 cycles for zenith satellites, with a linear increase to ±0.25 cycles for satellites at the horizon. These thresholds were found empirically.

Initially, we intended that our experimental design would employ the Cornell University GRID receiver and the prototype RF transmitter hardware pictured earlier to implement a spoofing attack against the GPS receiver and antenna array with a clear view of the sky, as follows:

1. Create a spoofing attack against the test receiver with the RASD firmware disabled and validate (using an oscilloscope) that the spoofer was able to drive the PPS generated by the target receiver away from truth.
2. Enable the RASD firmware and execute the same attack, this time validating that the receiver was able to detect the attack and raise an alarm.

Our 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!

In view of this, we elected 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 4. (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.)

We also moved the antenna array indoors and mounted it near to the retransmit antenna. In this environment the receiver was able to track the retransmitted signal without cycle slips when the range to the transmit antenna was less than approximately six meters. The indoor receiver was also able to decode all navigation data and calculate a position fix (the position of the rooftop antenna).

Figure 5 plots the delta phase profiles of seven satellite signals tracked during the indoor experiment. Consistent with the geometry, one may observe that all delta phases from the point source lie on top of each other.

The multi-receiver approach only requires that the vector between the two antennas is constant and known.

Having calculated the position fix and the LOS vectors, the RASD-enabled receiver immediately detected and flagged the retransmitted signals as coming from a "spoofer" not from the satellites themselves. Although this experiment did not detect an actual spoofer, the situation is sufficiently representative that the approach and software implementation were considered successful.

Conclusions

Antenna diversity — employing either multiple separate receivers or a multi-antenna single-oscillator receiver — can be used to defend against intentional GPS spoofing by greatly increasing the

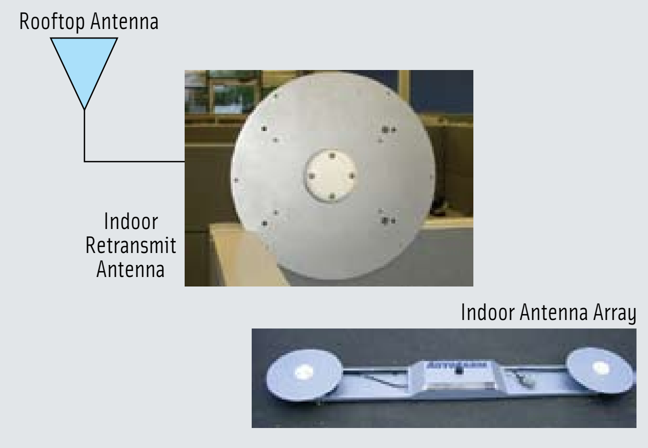

FIGURE 4 Indoor Experimental Setup

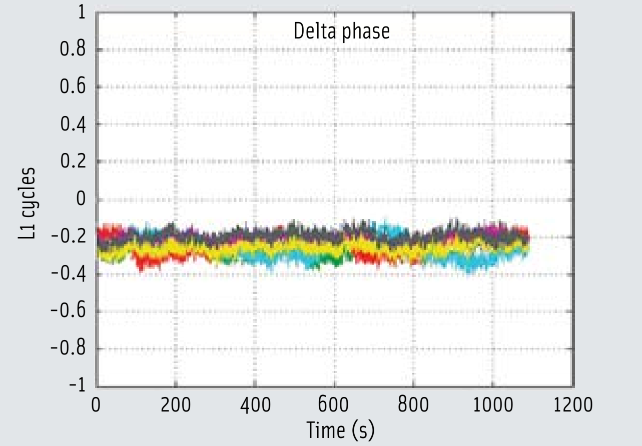

FIGURE 5 Experimental delta phase results with point (re)transmitter

technical difficulty required to mount a successful 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.

> **The technology to enable multi-antenna spoofing detection is readily available using any of the numerous GPS receivers that produce L1 carrier phase observables.**

If the GPS antennas of static or dynamic installations are further protected by physical security, it is possible to create a robust defense against even a sophisticated spoofing attack. In 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.

In the spoofing defense implemented here, a one-time survey of a fixed antenna array was sufficient to enable receiver autonomous spoofing detection. A practical but slightly less robust defense that does not depend on knowledge of the attitude of the multi-antenna array can also be implemented.

The technology to enable multi-antenna spoofing detection is readily available using any of the numerous GPS receivers that produce L1 carrier phase observables.

Acknowledgments

This article is based substantially on material in a paper first presented January 26, 2009, in Anaheim, California, at the International Technical Conference of the Institute of Navigation.

The authors would like to thank Novariant for the use of the AutoFarm roof array used for the experiment. Special thanks to Dennis Connor of Novariant for supporting this work. Additional thanks to William J. Bencze for RF hardware development support.

Manufacturers

The dual-antenna array used in the experiments described in this article was the AutoFarm antenna from **Novariant, Inc.**, Fremont, California, USA. The array's internal GPS receiver is based on the GP2015/GP2021 chipset, **Zarlink Semiconductor Inc.**, Ottawa, Ontario, Canada, and uses Novariant proprietary software. The spoofed handheld receivers were the eTrex from Garmin International, Olathe, Kansas, USA, and the iPhone from Apple Inc., Cupertino, California, USA.

Additional Resources

[1] Baldini, G., and J. Hofherr, "IPSC Projects based on Satellite Navigation Systems," 1st MENTORE Event, Institute for the Protection and Security of the Citizen, European Commission Joint Research Center, Ispra, Italy, November 25, 2008

[2] Commission Regulation (EC) No 2244/2003, "Laying down detailed provisions regarding satellite-based Vessel Monitoring Systems," *Official Journal of the European Union,* L 333/17, December 12, 2003, Brussels, Belgium

[3] Humphreys, T. E., and B. M. Ledvina, M. L. Psiaki, B. W. O' Hanlon, and P. M. Kintner, Jr., "Assessing the Spoofing Threat: Development of a Portable GPS Civilian Spoofer," *Proceedings of ION GNSS 2008,* Institute of Navigation, Savanna, Georgia, USA, 2008

[4] Key, E. L., "Techniques to Counter GPS Spoofing," internal memorandum, The MITRE Corporation, Bedord, Massachusetts, USA, February 1995

[5] "United States Positioning, Navigation, and Timing Interference Detection and Mitigation Plan Summary," U.S. Department of Homeland Security, Washington, D.C., USA, April 2008

[6] *"Vulnerability Assessment Of the Transportation Infrastructure Relying on the Global Positioning System,"* Technical Report, U.S. Department of Transportation, John A. Volpe National Transportation Systems Center, Cambridge, Massachusetts, USA, 2001

Authors

Paul Montgomery is the principal engineer at Novariant, Inc. He received a Ph.D in aeronautics and astronautics from Stanford University and was a founding member of Novariant, where his accomplishments have included adapting integrity beacon landing system (IBLS) technology to the Outrider TUAV (tactical unmanned aerial vehicle) 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.

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. IG