GNSS Spoofing Detection using Two-Antenna Differential Carrier Phase

Mark L. Psiaki, Brady W. O'Hanlon, and Steven P. Powell,
Cornell University, Ithaca, NY 14853-7501, U.S.A.

Jahshan A. Bhatti, Kyle D. Wesson, and Todd E. Humphreys,
The University of Texas at Austin, Austin, Texas 78712-0235, U.S.A.

Andrew Schofield,
Sea ID, 06320, Cap d'Ail, France

BIOGRAPHIES

Mark L. Psiaki is a Professor of Mechanical and Aerospace Engineering. He received a B.A. in Physics and M.A. and Ph.D. degrees in Mechanical and Aerospace Engineering from Princeton University. His research interests are in the areas of GNSS technology and applications, spacecraft attitude and orbit determination, and general estimation, filtering, and detection.

Brady W. O'Hanlon is a graduate student in the School of Electrical and Computer Engineering. He received a B.S. in Electrical and Computer Engineering from Cornell University. His interests are in the areas of GNSS technology and applications, GNSS security, and space weather.

Steven P. Powell is a Senior Engineer with the GPS and Ionospheric Studies Research Group in the Department of Electrical and Computer Engineering at Cornell University. He has M.S. and B.S. degrees in Electrical Engineering from Cornell University. He has been involved with the design, fabrication, testing, and launch activities of many scientific experiments that have flown on high altitude balloons, sounding rockets, and small satellites. He has designed ground-based and space-based custom GPS receiving systems primarily for scientific applications.

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

Kyle D. Wesson received his M.S. and Ph.D. in the Department of Electrical and Computer Engineering at The University of Texas at Austin. He received his B.S. in Electrical and Computer Engineering from Cornell University. He is a member of the UT Radionavigation Laboratory and the Wireless Networking and Communications Group. His research interests include GNSS security and interference mitigation.

Todd E. Humphreys is an assistant professor in the department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin, and Director of the UT Radionavigation Laboratory. He received a B.S. and M.S. in Electrical and Computer Engineering from Utah State University and a Ph.D. in Aerospace Engineering from Cornell University. He specializes in applying optimal estimation and signal processing techniques to problems in radionavigation. His recent focus is on radionavigation robustness and security.

Andrew Schofield is a career Yacht Captain. After completing his degree in Applied Biology and working in the bio-science industry for a year, he left all that behind in 1991 and found a deck hand's job on a sailing yacht in the Caribbean. Since then he has worked on various yachts in various locations. He has been Captain of the White Rose of Drachs since launch in June 2004. He is President of the Professional Yachting Association, the large yacht professional body, and focuses on the training and certification of crew. In his time at sea GPS has transformed navigation. He feels that the relevance of the work done to detect GPS spoofing cannot be overstated with regard to the safety of life at sea, and he is delighted to have facilitated the voyage during which spoofing detection was proven.

ABSTRACT

A method is developed to detect GNSS spoofing by processing beat carrier-phase measurements from a pair of antennas in a CDGPS-type calculation. This system detects spoofing attacks that are resistant to standard RAIM technique, and it can sense an attack in a fraction of a second without external aiding. The signal-in-space properties used to detect spoofing are the relationships of the signal arrival directions to the vector that points from one antenna to the other. In the un-spoofed case, there are a multiplicity of relationships between the inter-antenna vector and the arrival directions of the multiple signals, which results in a quantifiable multiplicity of carrier-phase single-differences between the antennas. In the spoofed case, there is a single direction of arrival, assuming a single spoofer transmission antenna, and the carrier phase single-differences are identical for all channels, up to an integer cycle ambiguity. A real-time implementation of this detection method has been developed, and it has been tested against live-signal spoofing attacks aboard a superyacht that was cruising around Italy en route from Monaco to Venice. The prototype system demonstrated an ability to detect spoofing attacks in a fraction of a second, though lags in the system's signal processing lengthened the detection delay to as much as 6 seconds. The system experienced challenges during the initial phase of a spoofing attack if the spoofer power was not much greater than that of the true signal. The true and spoofed signals interfere in a beating pattern in this case, making the composite signal harder to track and harder to classify as being either spoofed or non-spoofed. After the spoofer drags the victim receiver off to an erroneous position or timing fix, the beating subsides, and the new spoofing detection system performs well.

INTRODUCTION

Concerns about spoofing of open-service GNSS signals inspired early work on simple RAIM methods based on the consistency of the navigation solution [1]. Work on new classes of defense techniques began in earnest after the demonstration of a powerful spoofer that is undetectable by simple pseudorange-based RAIM methods [2,3]. There has been a sense of urgency to solve the spoofing problem since the Iranians captured a highly classified U.S. drone in Dec. 2011 and made unsubstantiated claims to have spoofed its GPS [4]. A pair of dramatic field demonstrations of the spoofer of Refs. 2 and 3 have further heightened interest in the GNSS spoofing detection problem. One test involved deception of a small UAV, causing it to dive towards the ground when it had been in hover mode [5,6]. Another test sent a superyacht off course without raising any alarms on its bridge [7,8].

Many contributions have been made towards the development and evaluation of methods that can detect attacks by the new class of spoofers [9-54]. One class of detection methods uses encrypted signals, their known relationships to the open-service signals, and after-the-fact availability of encryption information in order to detect spoofing [11,13,18,42,44,46,47,50]. Such techniques require a high-bandwidth communication link between the potential victim of a spoofing attack and a trusted source of after-the-fact encryption information [46]. They may involve significant latency between a spoofing attack and a detection.

Another class of methods uses advanced RAIM-type techniques [14,16,19,22,26,28,32,37,43,45,49,51,52,54]. Instead of considering only pseudorange consistency, advanced RAIM techniques examine additional signal characteristics such as absolute power levels, distortion of the PRN code correlation function along the early/late axis, the possible existence of multiple distinct correlation peaks in signal-acquisition-type calculations, and other signal or receiver characteristics. Such methods are relatively simple to implement because they do not require much additional hardware, if any, but some of these strategies can have trouble distinguishing between multipath and spoofing [19] or between jamming and spoofing.

A third class of methods proposes the addition of Navigation Message Authentication bits [20,30,35]. These are encrypted parts of the low-rate navigation data message. Such techniques require modification of the navigation data message and can allow long latencies between the onset of a spoofing attack and its detection [55].

A fourth class of methods exploits the differing signal-in-space geometry of spoofed signals in comparison to true GNSS signals. All spoofed signals typically arrive from the same direction, but true signals arrive from a multiplicity of directions. Some of these methods use receiver antenna motion to achieve direction-of-arrival sensitivity [12,15,21,23,36,39,41,53]. Others use an array of two or more receiver antennas to exploit the difference of the arrival direction distributions between spoofed and non-spoofed situations [9,10,24,25,27,31,33,34,38,40,48]. The most powerful of these detection strategies exploit models of the effects on carrier phase data of antenna motion or antenna array geometry [9,27,33,36,40,41]. This knowledge may be partial because an unknown antenna array attitude may need to be determined as part of the spoofing detection calculation. Such calculations amount to spoofing detection augmentations of standard CDGPS attitude determination. Their detection power derives from the high degree of accuracy with which a typical GNSS receiver can measure beat carrier phase.

The present effort is a follow-on to the moving-antenna/carrier-phase-based spoofing detection work reported in Refs. 36 and 41. One goal of this effort has been to remove the necessity for actual moving parts by using two antennas and processing their carrier-phase data in a way that is analogous to the moving-antenna scheme of Refs. 36 and 41. A second goal has been to achieve real-time operation. The prototype moving-antenna system reported in Refs. 36 and 41 used post-processing and completed its spoofing detection calculations days or weeks after the recording of wide-band RF data during live-signal attacks. In theory, nothing prevents the methods of Refs. 36 and 41 from being implemented in real-time, and one goal of this study has been to provide a practical demonstration of this fact. A third goal has been to test this system against actual live-signal spoofing attacks in order to prove its real-time capabilities and evaluate its performance during the two phases of an attack, the initial signal capture and the post-capture drag-off to erroneous position and timing fixes.

This paper makes five contributions to the art of spoofing detection. The first is to define the architecture of a two-antenna system that can perform spoofing detection without employing any moving parts. The second contribution is the development of a mathematical model of the two-antenna system and a corresponding spoofing detection hypothesis test. This test involves optimal estimation calculations that are modifications of those used in Refs. 36 and 41. The third contribution is an explanation of how single-differenced beat carrier phase can be determined from a switched-antenna version of the two-antenna system, one that uses an RF switch between the two antennas, a single RF front-end, and a single receiver channel per tracked signal. Typically one would need multiple RF front-ends and multiple receiver channels per satellite, i.e., one for each antenna/satellite combination. The fourth contribution is an evaluation of the real-time performance of this system during live-signal spoofing attacks which were conducted aboard the same superyacht that was used in the studies described in Refs. 7 and 8. The fifth contribution is a set of recommendations for future improvements and studies along with an explanation of how recorded data from the live-signal attacks can be used to evaluate enhanced spoofing detection systems.

The remainder of this paper is divided into 5 main sections plus conclusions. Section II describes the hardware and functioning of 2 possible versions of the two-antenna spoofing detection system. Section III develops a spoofing detection hypothesis test that operates on the single-differenced carrier phases from the two-antenna system. Section IV explains how a special Phase-Lock Loop (PLL) can be developed to track any given GNSS signal from a switched-antenna version of the system. It returns the single-differenced carrier phase as part of its signal processing outputs. Section V describes the tests that have been performed on this system, including preliminary off-line tests and live-signal tests aboard the yacht. Section VI proposes ways to enhance the basic spoofing detection system in order to address anomalies noted in Section V. It also explains how the enhancements can be tested without the need for a new live-signal test campaign. Section VII summarizes this paper's results and gives its conclusions.

II. ARCHITECTURE OF TWO-ANTENNA SPOOFING DETECTION SYSTEM

The spoofing detection system consists of two GNSS patch antennas, GPS receiver hardware and software, and spoofing detection signal processing hardware and software. Two versions of this system are depicted in Fig. 1. The left-hand version connects its two patch antennas to an RF switch. The single analog RF output of the switch is input to a GNSS receiver that is standard in all respects, except for two features. First, it controls the RF switch or, at least, has access to the switching times. Second, it employs a specialized PLL that can track the beat carrier phase of a given signal through the phase jumps that occur at the switching times. The right-hand version connects each of its two antennas to an independent GPS receiver, likely connected to a common reference oscillator.

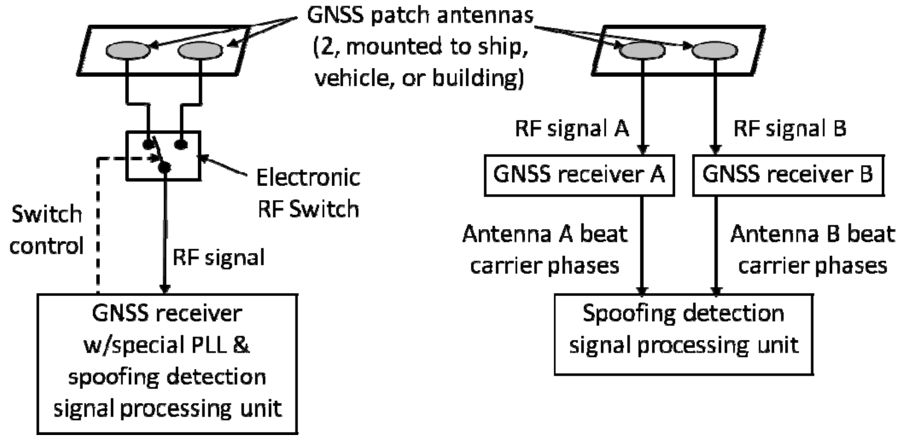

Fig. 1. Two configurations of the two-antenna spoofing detection system, the RF-switched-signal/single-receiver configuration (left) and the two-receiver configuration (right).

The last element of each system is a spoofing detection signal processing unit. Its inputs are the single-differenced beat carrier phases of all tracked signals, with differences taken between the two antennas. In the switched antenna system, each signal difference is deduced by the specialized PLL that tracks beat carrier phase through the antenna switching times. In the two-receiver system, the single-differences are calculated explicitly from each receiver's beat carrier phase observables. The calculations of the spoofing detection signal processing unit are the subject of Section III of this paper. Section IV gives an overview of how the special PLL for the switched-antenna system would work.

Except for the final spoofing detection unit, the two-receiver system on the right-hand side of Fig. 1 is already available commercially. Typical applications are CDGPS-based attitude/heading determination. Thus, this is the easiest version to implement. It is the version that is tested in Section V of this paper.

Although it cannot be implemented with off-the-shelf hardware and signal processing algorithms, the switched-antenna system on the left-hand side of Fig. 1 has some advantages. First, it reduces the number of RF front-ends from two to one. Second, it cuts the number of receiver tracking channels in half. These reductions come at the additional cost of adding the RF switch and the switching control/information connection to the receiver.

The switched-antenna system can be viewed as being a sort of moving-antenna system, like that of Refs. 36 and 41. Instead of using actual physical motion of the antenna, the switch creates step changes in the effective antenna phase center location that are the square-wave equivalent of the oscillatory physical motions used in the system of Refs. 36 and 41.

This system could include more than two antennas, as in Refs. 27, 33, and 40. A multi-antenna system could have a dedicated RF front-end and a dedicated set of receiver channels for each antenna, as in the two-antenna system on the right-hand side of Fig. 1. Alternatively, a multi-antenna system could include an RF switch that can switch between any one of the multiple antennas at the command of the receiver. The latter design would entail a slight modification to the specialized PLL in order to track multiple independent phase jumps for the independent antenna switches.

The principles used to detect spoofing can be understood by considering and comparing the signal-in-space and antenna geometries shown in Fig. 2 and Fig. 3. Fig. 2 shows the two-antenna system and three GNSS satellites for a typical non-spoofed case. Fig. 3 shows a spoofed case. The salient difference is that the different GNSS signals arrive from different directions for the non-spoofed case of Fig. 2, namely $\hat{\boldsymbol{r}}^{j-1}$, $\hat{\boldsymbol{r}}^{j}$, and $\hat{\boldsymbol{r}}^{j+1}$. They all arrive from the same direction, the direction of the spoofer $\hat{\boldsymbol{r}}^{sp}$, for the spoofed case of Fig. 3. For spoofing detection purposes, the important geometric feature is the projection of each direction of arrival onto the known separation vector between the two antennas, $\boldsymbol{b}_{BA}$. This projection has a direct effect on the beat carrier phase difference between the two antennas. In the non-spoofed case, this effect will vary between the different received signals in ways consistent with the attitude of the $\boldsymbol{b}_{BA}$ vector. In the spoofed case, all of these carrier phase differences will be identical. The spoofing detection algorithm decides between two hypotheses about the carrier-phase differences, one conjecturing a diversity consistent with authentic signals and the other conjecturing the sameness that is characteristic of spoofed signals.

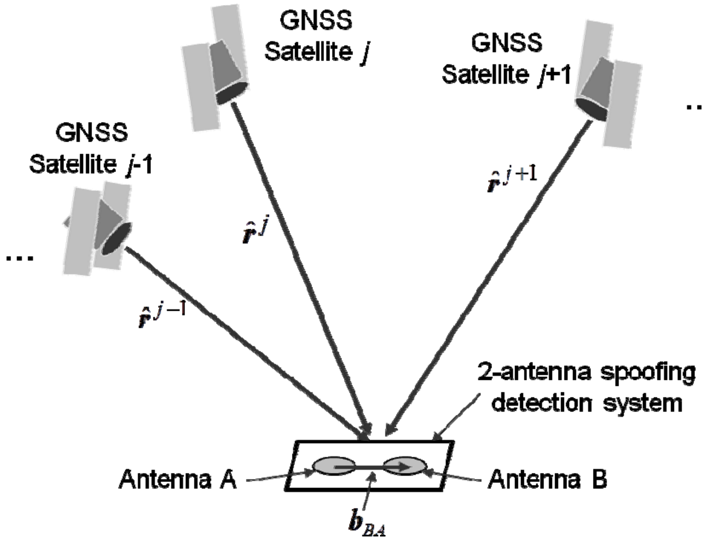
GNSS Satellite *j*
GNSS Satellite *j*+1
...
GNSS Satellite *j*-1
...
$\hat{\boldsymbol{r}}^{j+1}$
$\hat{\boldsymbol{r}}^{j}$
$\hat{\boldsymbol{r}}^{j-1}$
2-antenna spoofing detection system
Antenna A
Antenna B
$\boldsymbol{b}_{BA}$

Fig. 2. Geometry of two-antenna spoofing detection system and GNSS satellites for non-spoofed case.

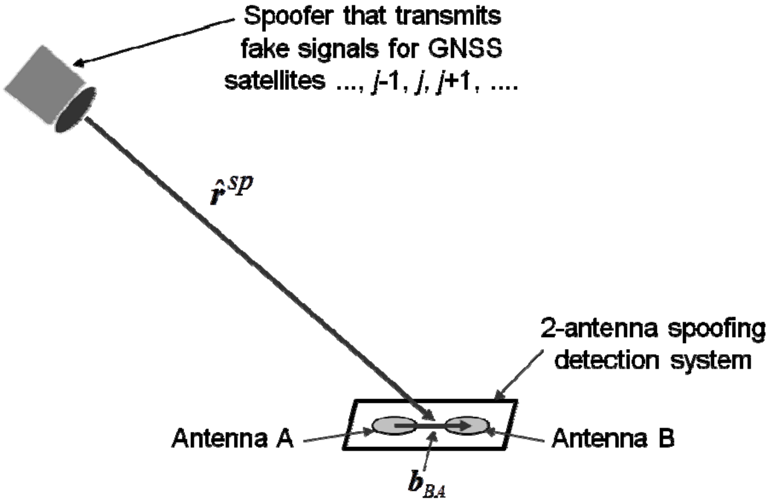
Spoofer that transmits fake signals for GNSS satellites ..., *j*-1, *j*, *j*+1, ...
$\hat{\boldsymbol{r}}^{sp}$
2-antenna spoofing detection system
Antenna A
Antenna B
$\boldsymbol{b}_{BA}$

Fig. 3. Spoofed-case geometry of two-antenna spoofing detection system and GNSS spoofer.

III. TWO-ANTENNA SPOOFING DETECTION HYPOTHESIS TEST

This section develops the spoofing detection hypothesis test. It starts by presenting the non-spoofed and spoofed signal models that form the basis of the test. Next, it develops optimal estimation algorithms that fit the observed differential beat carrier phases to either of the two models. Last, it shows how these two estimates and their associated fit error costs can be used to develop a sensible spoofing detection hypothesis test.

A. Mathematical Models of Non-Spoofed and Spoofed Cases

The non-spoofed single-differenced carrier phase model takes the form:

$$\begin{aligned}\Delta\phi_{BA}^{j} &= \phi_{B}^{j} - \phi_{A}^{j} \\ &= -\frac{2\pi}{\lambda}(\hat{\boldsymbol{r}}^{j})^{\mathrm{T}} A^{\mathrm{T}}\boldsymbol{b}_{BA} \\ &\qquad + \beta + 2\pi\Delta N_{BA}^{j} + n_{mpBA}^{j} + n_{rcvrBA}^{j} \\ &= -\frac{2\pi\rho_{BA}}{\lambda}(\hat{\boldsymbol{r}}^{j})^{\mathrm{T}}\hat{\boldsymbol{r}}_{BA} \\ &\qquad + \beta + 2\pi\Delta N_{BA}^{j} + n_{mpBA}^{j} + n_{rcvrBA}^{j} \qquad (1)\end{aligned}$$

where ϕ_{A}^{j} and ϕ_{B}^{j} are the (negative) beat carrier phases of the signal from GNSS satellite j received at, respectively, Antennas A and B. These phases are termed negative because they are the time integrals of the carrier Doppler shifts and have the opposite sign of the usual beat carrier phase definition in the GNSS literature.

The single-differenced carrier phase observable $\Delta\phi_{BA}^{j}$ on the left-hand side of Eq. (1) is modeled on the right hand sides of the last two lines as the sum of a leading geometric term, the line-bias-plus-fractional-differential-phase term β, the single-differenced integer-ambiguity term $2\pi\Delta N_{BA}^{j}$, the single-differenced multipath noise term n_{mpBA}^{j}, and the single-differenced receiver thermal noise term n_{rcvrBA}^{j}.

The geometric term includes the known unit direction vector from GNSS satellite j to the centroid of the 2 antennas as defined in reference coordinates, $\hat{\boldsymbol{r}}^{j}$, the known vector from Antenna A's phase center to Antenna B's phase center as defined in local antenna body coordinates, $\boldsymbol{b}_{BA}$, the unknown 3x3 orthonormal direction cosines matrix that transforms from reference coordinates to body coordinates, A, and the nominal carrier wavelength, λ. The last line of Eq. (1) eliminates the known antenna baseline vector $\boldsymbol{b}_{BA}$ and the unknown direction cosines matrix A in favor of the known antenna baseline length $\rho_{BA} = (\boldsymbol{b}_{BA}^{\mathrm{T}}\boldsymbol{b}_{BA})^{0.5}$ and the unknown unit direction vector from Antenna A to Antenna B given in reference coordinates: $\hat{\boldsymbol{r}}_{BA} = A^{\mathrm{T}}\boldsymbol{b}_{BA}/\rho_{BA}$.

The spoofed single-differenced carrier phase model is:

$$\begin{aligned}\Delta\phi_{BA}^{j} &= -\frac{2\pi\rho_{BA}}{\lambda}(\hat{\boldsymbol{r}}^{sp})^{\mathrm{T}}\hat{\boldsymbol{r}}_{BA} \\ &\qquad + \beta + 2\pi\Delta N_{BA}^{j} + n_{mpBA}^{sp} + n_{rcvrBA}^{j} \\ &= \beta_{sp} + 2\pi\Delta N_{BA}^{j} + n_{rcvrBA}^{j} \qquad (2)\end{aligned}$$

There are two differences between this model and the non-spoofed model. One is the replacement of the satellite-to-antennas direction vector $\hat{\boldsymbol{r}}^{j}$ by the spoofer-to-antennas direction vector $\hat{\boldsymbol{r}}^{sp}$ in the geometric term on the right-hand side of the first line. The other is the re-designation of the multipath noise term to include the superscript $()^{sp}$ instead of $()^{j}$ in order to indicate that the multipath error is identical for all spoofed signals.

One can lump the geometric term, the original line-bias term, and the multipath term on the first line of Eq. (2) into a new re-defined line-bias term in order to derive the second line of Eq. (2):

$$\beta_{sp} = -\frac{2\pi\rho_{BA}}{\lambda}(\hat{\boldsymbol{r}}^{sp})^{\mathrm{T}}\hat{\boldsymbol{r}}_{BA} + \beta + n_{mpBA}^{sp} \qquad (3)$$

This lumping together of unknowns is reasonable because none of the terms on the right-hand side of Eq. (3) depends on the satellite identifier superscript j. This aggregation of terms is important for 2 reasons. First, it reduces the number of unknown parameters under the spoofing hypothesis. Second, the total noise standard deviation is reduced due to the elimination of multipath as a source of differential noise between different signals.

B. Optimal Estimation of Unknown Parameters for Non-Spoofed and Spoofed Cases

As in Refs. 36 and 41, this paper's spoofing detection statistic is based on a comparison of how well the non-spoofed and spoofed signal models fit the actual data. Both models contain unknown parameters. Therefore, optimal estimation problems are solved for each model in order to assess how well it fits the data.

Suppose that there are L single-differenced carrier phase signals available for spoofing detection. Then the following optimal estimation problem is used to assess the fit of the non-spoofed model to the data

find: $\hat{\boldsymbol{r}}_{BA}$, β, and ΔN_{BA}^{1}, …, ΔN_{BA}^{L} (4a)

to minimize: $J_{nonsp}(\hat{\boldsymbol{r}}_{BA}, \beta, \Delta N_{BA}^{1}, \ldots, \Delta N_{BA}^{L}) =$

$$\frac{1}{2}\sum_{j=1}^{L}\frac{[\Delta\phi_{BA}^{j}+\frac{2\pi\rho_{BA}}{\lambda}(\hat{\boldsymbol{r}}^{j})^{\mathrm{T}}\hat{\boldsymbol{r}}_{BA}-\beta-2\pi\Delta N_{BA}^{j}]^{2}}{\sigma_{mp}^{2}+(\sigma_{rcvr}^{j})^{2}} \tag{4b}$$

subject to: $(\hat{\boldsymbol{r}}_{BA})^{\mathrm{T}}\hat{\boldsymbol{r}}_{BA}=1$ (4c)

$\Delta N_{BA}^{1}=0$ (4d)

ΔN_{BA}^{j} integer-valued for j = 2, …, L (4e)

where σ_{mp} is the standard deviation of the multipath noise n_{mpBA}^{j} and σ_{rcvr}^{j} is the standard deviation of the receiver thermal noise n_{rcvrBA}^{j}. The two noise terms are assumed to be zero-mean Gaussian noise. The thermal noise standard deviation can vary between signals due to their differing carrier-to-noise ratios. The multipath noise, however, is assumed to have a single standard deviation for all signals. The constraint in Eq. (4c) enforces the unit normalization of the $\hat{\boldsymbol{r}}_{BA}$ direction vector. The Eq.-(4d) constraint of zero integer ambiguity for signal j = 1 serves to lump its ambiguity into the line bias term β, thereby preserving system observability.

The cost function in Eq. (4b) equals the negative natural logarithm of the probability density function of the single-differenced carrier phase data modeled in Eq. (1), except that the natural logarithm of the probability density's normalization constant has been omitted. The associated probability density function is conditioned on the unknown parameters given in the cost function's argument list. Therefore, the optimal estimation problem in Eqs. (4a)-(4e) is a maximum likelihood estimation problem.

The maximum-likelihood optimal estimation problem for the spoofed case can be posed in the following form:

find: β_{sp} and ΔN_{BA}^{1}, …, ΔN_{BA}^{L} (5a)

to minimize: $J_{sp}(\beta_{sp},\Delta N_{BA}^{1},\ldots,\Delta N_{BA}^{L})=$

$$\frac{1}{2}\sum_{j=1}^{L}\frac{[\Delta\phi_{BA}^{j}-\beta_{sp}-2\pi\Delta N_{BA}^{j}]^{2}}{(\sigma_{rcvr}^{j})^{2}} \tag{5b}$$

subject to: $\Delta N_{BA}^{1}=0$ (5c)

ΔN_{BA}^{j} integer-valued for j = 2, …, L (5d)

This optimal estimation problem differs from that in Eqs. (4a)-(4e) in two significant ways. First, there is no unit direction vector $\hat{\boldsymbol{r}}_{BA}$ to estimate and no corresponding unit-normalization constraint. Second, the normalizing variance in the denominator of the summand is smaller due to the absence of the multipath variance term. The correctness of using this lowered variance has been deduced from the identical nature of the multipath errors in the spoofed case, and it has been verified experimentally.

Algorithm for Solving Spoofed-Case Estimation Problem. Solution of the spoofed-case estimation problem in Eqs. (5a)-(5d) is straightforward and can be accomplished analytically. For a given guess of the modified line bias β_{sp}, the optimal integer ambiguities can be computed by rounding:

$$\Delta N_{BAopt}^{1}(\beta_{sp})=0 \tag{6a}$$

$$\Delta N_{BAopt}^{j}(\beta_{sp})=round\left(\frac{\Delta\phi_{BA}^{j}-\beta_{sp}}{2\pi}\right) \text{ for } j = 2, \ldots, L \tag{6b}$$

where *round*() is the usual function that rounds to the nearest integer. The formula in Eq. (6b) can be used to reduce the problem in Eqs. (5a)-(5d) to a one-dimensional unconstrained search in the real-valued variable β_{sp} in order to minimize the following modified cost function

$$\tilde{J}_{sp}(\beta_{sp})=\frac{1}{2}\sum_{j=1}^{L}\frac{[\Delta\phi_{BA}^{j}-\beta_{sp}-2\pi\Delta N_{BAopt}^{j}(\beta_{sp})]^{2}}{(\sigma_{rcvr}^{j})^{2}} \tag{7}$$

The sum of the 2nd through Lth cost terms in Eq. (7) is periodic in β_{sp} with period equal to 2π as a result of the rounding operation in Eq. (6b). Therefore, one can easily show that the optimal value of β_{sp} must lie in the range $\Delta\phi_{BA}^{1}-\pi\le\beta_{sp}\le\Delta\phi_{BA}^{1}+\pi$ because the first term in the cost sum in Eq. (7) is smaller in that β_{sp} range than in any range of the form $\Delta\phi_{BA}^{1}+(2N–1)\pi\le\beta_{sp}\le\Delta\phi_{BA}^{1}+(2N+1)\pi$ for non-zero integer N.

The optimal β_{sp} can be determined by noting that the cost function in Eq. (7) is continuous and piecewise quadratic in β_{sp}. The node points at which it changes from one piecewise quadratic model to another in the range $\Delta\phi_{BA}^{1}-\pi\le\beta_{sp}\le\Delta\phi_{BA}^{1}+\pi$ are

$$\beta_{spnode}^{j}=\Delta\phi_{BA}^{j}+2\pi\, ceil\left(\frac{\Delta\phi_{BA}^{1}-\Delta\phi_{BA}^{j}}{2\pi}\right)-\pi \quad \text{for } j = 2, \ldots, L \tag{8}$$

where the *ceil*() function produces the integer nearest the input argument and no less than it, i.e., no closer to $-\infty$. Between each neighboring pair of node points, say β_{spnode}^{i} and β_{spnode}^{k}, all of the values of $\Delta N_{BAopt}^{j}(\beta_{sp})$ for j = 1, …, L remain constant, and it is straightforward to optimize the resulting quadratic function in β_{sp} over this interval. There are L such intervals in the range $\Delta\phi_{BA}^{1}-\pi\le\beta_{sp}\le\Delta\phi_{BA}^{1}+\pi$, and brute-force consideration of all these intervals yields the global optimal solution to the problem in Eqs. (5a)-(5d) in a small number of computations, a number that varies linearly with the number of signals L.

Algorithm for Solving Non-Spoofed-Case Estimation Problem. Solution of the non-spoofed-case estimation problem in Eqs. (4a)-(4e) is more difficult. Algorithms from the general CDGPS attitude determination literature could be adapted to this purpose, such as those of Refs. 56 and 57. A new approach has been developed here that is appropriate for a short baseline and this seemingly under-determined problem; the number of single-differenced carrier phases, L, is lower than the difference between the number of unknowns and the number of equality constraints, L+2. In reality, the constraints that the ambiguities be integer-valued makes this problem well determined for $L \geq 3$.

The algorithm finds the global minimum solution to the problem in Eqs. (4a)-(4e) by calling a local minimizer multiple times. The local minimize starts with a guess of the unit direction vector between the two antennas in reference coordinates, $\hat{\boldsymbol{r}}_{BAguess}$. It iterates from this guess to find a locally optimum solution $\hat{\boldsymbol{r}}_{BAopt}$, β_{opt}, and ΔN_{BAopt}^{1}, …, ΔN_{BAopt}^{L}. The global solution algorithm re-starts the local minimize from a set of $\hat{\boldsymbol{r}}_{BAguess}$ values. The global minimizer chooses its set of $\hat{\boldsymbol{r}}_{BAguess}$ values so that every $\hat{\boldsymbol{r}}_{BA}$ value on the unit sphere satisfies the inequality

$$\hat{\boldsymbol{r}}_{BAguess}^{\mathrm{T}} \hat{\boldsymbol{r}}_{BA} \geq \cos\left(\frac{\lambda}{4\rho_{BA}} \right) \tag{9}$$

for at least one of the $\hat{\boldsymbol{r}}_{BAguess}$ vectors. This constraint implies that the attitude term in Eq. (4b) will vary by no more than $\pi/2$ radians when iterating from the nearest initial guess to the global optimal solution. In that case, none of the integer ambiguities should change. For a problem with long baselines, i.e., $\rho_{BA} >> \lambda$, such an approach might entail a very large number of initial guesses $\hat{\boldsymbol{r}}_{BAguess}$, and it might be more efficient to adapt one of the algorithms of Refs. 56 or 57. For the system tested in Section V, $\rho_{BA} = 0.74\lambda$, and this algorithm is reasonably efficient.

The local minimization algorithm is a heuristic procedure that alternates between optimizing two overlapping subsets of the unknowns in Eq. (4a) while leaving the other unknowns at their current guesses. One optimized subset is {β, ΔN_{BA}^{1}, …, ΔN_{BA}^{L}}, and the other is {$\hat{\boldsymbol{r}}_{BA}$, β,}. Optimization of {β, ΔN_{BA}^{1}, …, ΔN_{BA}^{L}} for a fixed guess of $\hat{\boldsymbol{r}}_{BA}$ can be accomplished using the exact same algorithm as has been defined above to solve the problem in Eqs. (5a)-(5d), except each $\Delta\phi_{BA}^{j}$ is replaced by $\Delta\phi_{BA}^{j} + (2\pi\rho_{BA}/\lambda)(\hat{\boldsymbol{r}}^{j})^{\mathrm{T}}\hat{\boldsymbol{r}}_{BA}$ for j = 1, …, L in order to account for the effects of the current guess of $\hat{\boldsymbol{r}}_{BA}$.

Optimization of {$\hat{\boldsymbol{r}}_{BA}$, β,} for a fixed guess of ΔN_{BA}^{1}, …, ΔN_{BA}^{L} can be accomplished by using an adaptation of algorithms defined in Ref. 41. This procedure starts by defining the following system of linear equations, which may or may not be over-determined:

$$\begin{bmatrix} \left(\dfrac{\Delta\phi_{BA}^{1} - 2\pi\Delta N_{BA}^{1}}{\sqrt{(\sigma_{rcvr}^{1})^{2} + \sigma_{mp}^{2}}} \right) \\ \left(\dfrac{\Delta\phi_{BA}^{2} - 2\pi\Delta N_{BA}^{2}}{\sqrt{(\sigma_{rcvr}^{2})^{2} + \sigma_{mp}^{2}}} \right) \\ \vdots \\ \left(\dfrac{\Delta\phi_{BA}^{L} - 2\pi\Delta N_{BA}^{L}}{\sqrt{(\sigma_{rcvr}^{L})^{2} + \sigma_{mp}^{2}}} \right) \end{bmatrix} = \begin{bmatrix} \left(\dfrac{1}{\sqrt{(\sigma_{rcvr}^{1})^{2} + \sigma_{mp}^{2}}} \right) & \left(\dfrac{-2\pi\rho_{BA}(\hat{\boldsymbol{r}}^{1})^{\mathrm{T}}}{\lambda\sqrt{(\sigma_{rcvr}^{1})^{2} + \sigma_{mp}^{2}}} \right) \\ \left(\dfrac{1}{\sqrt{(\sigma_{rcvr}^{2})^{2} + \sigma_{mp}^{2}}} \right) & \left(\dfrac{-2\pi\rho_{BA}(\hat{\boldsymbol{r}}^{2})^{\mathrm{T}}}{\lambda\sqrt{(\sigma_{rcvr}^{2})^{2} + \sigma_{mp}^{2}}} \right) \\ \vdots & \vdots \\ \left(\dfrac{1}{\sqrt{(\sigma_{rcvr}^{L})^{2} + \sigma_{mp}^{2}}} \right) & \left(\dfrac{-2\pi\rho_{BA}(\hat{\boldsymbol{r}}^{L})^{\mathrm{T}}}{\lambda\sqrt{(\sigma_{rcvr}^{L})^{2} + \sigma_{mp}^{2}}} \right) \end{bmatrix} \begin{bmatrix} \beta \\ \hat{\boldsymbol{r}}_{BA} \end{bmatrix} \tag{10}$$

or

$$\boldsymbol{y} = H \begin{bmatrix} \beta \\ \hat{\boldsymbol{r}}_{BA} \end{bmatrix} \tag{11}$$

where the L-by-1 vector $\boldsymbol{y}$ and the L-by-4 matrix H in Eq. (11) are defined by the correspondence between Eqs. (10) and (11). Note how the current guesses of the integer ambiguities ΔN_{BA}^{1}, …, ΔN_{BA}^{L} are incorporated into the definition of the $\boldsymbol{y}$ vector on the right-hand side of Eq. (10).

The {$\hat{\boldsymbol{r}}_{BA}$, β,} optimization minimizes the sum of the squared errors in Eq. (11) subject to the $\hat{\boldsymbol{r}}_{BA}$ unit normalization constraint in Eq. (4c). It starts with an orthonormal/upper-triangular (QR) factorization [58] of the H matrix:

$$Q \begin{bmatrix} R_{\beta\beta} & R_{\beta r} \\ 0 & R_{rr} \\ 0 & 0 \end{bmatrix} = H \tag{12}$$

The input to the QR factorization is the H matrix on the right-hand side of Eq. (12), and the outputs are the L-by-L orthonormal matrix Q and the upper-triangular factor that contains the scalar $R_{\beta\beta}$, the 1-by-3 matrix $R_{\beta r}$, and the 3-by-3 upper-triangular matrix R_{rr}. The Q matrix is also used to transform the $\boldsymbol{y}$ vector to yield

$$\begin{bmatrix} z_\beta \\ \boldsymbol{z}_r \\ \boldsymbol{z}_{resid} \end{bmatrix} = Q^{\mathrm{T}} \boldsymbol{y} \tag{13}$$

where z_β is a scalar, $\boldsymbol{z}_r$ is a 3-by-1 vector, and $\boldsymbol{z}_{resid}$ is an (L-4)-by-1 vector. Note that the use of this QR factorization requires $L \geq 4$ in order for this algorithm to be implementable.

Given the transformations in Eqs. (12) and (13), the optimal value of $\hat{\boldsymbol{r}}_{BA}$ is the value that minimizes the sum of the squared errors in the equation

$$\boldsymbol{z}_r = R_{rr} \hat{\boldsymbol{r}}_{BA} \tag{14}$$

subject to the unit normalization constraint on $\hat{\boldsymbol{r}}_{BA}$ found in Eq. (4c). This optimization problem can be solved in closed form by using a singular-value decomposition of R_{rr} and calculations found in Section V.A of Ref. 41. In order to implement the calculations of Ref. 41, one must use R_{rr} from Eq. (12) in place of the B matrix of Ref. 41, and one must use the $\boldsymbol{z}_r$ vector from Eq. (13) in place of the $[\, z_4^1, z_4^2, \ldots, z_4^L \,]^{\mathrm{T}}$ vector of Ref. 41. Given the minimizing $\hat{\boldsymbol{r}}_{BA}$ -- call it $\hat{\boldsymbol{r}}_{BAopt}$, the minimizing value of β is $\beta_{opt} = (z_\beta - R_{\beta r} \hat{\boldsymbol{r}}_{BAopt})/R_{\beta\beta}$.

The algorithm for finding a local minimum solution to the problem in Eqs. (4a)-(4e) starts with its initial guess of $\hat{\boldsymbol{r}}_{BA}$ and performs minimization with respect to β and $\Delta N_{BA}^1, \ldots, \Delta N_{BA}^L$, as discussed above. Next, it holds $\Delta N_{BA}^1, \ldots, \Delta N_{BA}^L$ fixed at their optimized values from this partial minimization, and it performs minimization with respect to $\hat{\boldsymbol{r}}_{BA}$ and β, as discussed above. This process then repeats with a new minimization with respect to β and $\Delta N_{BA}^1, \ldots, \Delta N_{BA}^L$ while holding $\hat{\boldsymbol{r}}_{BA}$ fixed at the optimal value determined from the most recent joint $\{\hat{\boldsymbol{r}}_{BA}, \beta\}$ minimization. The algorithm terminates when the optimal $\Delta N_{BA}^1, \ldots, \Delta N_{BA}^L$ values from a joint $\{\beta, \Delta N_{BA}^1, \ldots, \Delta N_{BA}^L\}$ minimization are the same as the values from the previous such minimization. This algorithm is guaranteed to arrive at a local minimum in $\{\hat{\boldsymbol{r}}_{BA}, \beta\}$ because every iteration will decrease the cost function in Eq. (4b) until the $\Delta N_{BA}^1, \ldots, \Delta N_{BA}^L$ values become static.

C. Spoofing Detection Hypothesis Test

The spoofing detection hypothesis test statistic is the difference of the optimal costs from the two optimization problems of the previous subsection:

$$\gamma = J_{sp}(\beta_{spopt}, \Delta N_{BAspopt}^1, \ldots, \Delta N_{BAspopt}^L) - J_{nonsp}(\hat{\boldsymbol{r}}_{BAopt}, \beta_{opt}, \Delta N_{BAnsopt}^1, \ldots, \Delta N_{BAnsopt}^L) \tag{15}$$

where $\{\beta_{spopt}, \Delta N_{BAspopt}^1, \ldots, \Delta N_{BAspopt}^L\}$ is the optimal solution to the spoofed-case estimation problem in Eqs. (5a)-(5d) and $\{\hat{\boldsymbol{r}}_{BAopt}, \beta_{opt}, \Delta N_{BAnsopt}^1, \ldots, \Delta N_{BAnsopt}^L\}$ is the optimal solution to the non-spoofed-case estimation problem in Eqs. (4a)-(4e).

This spoofing detection statistic is similar to those defined in Ref. 41 for various spoofing detection cases. It is analogous to a Neyman-Pearson test [59] because it is traceable to a ratio of the probability densities of the data given the two hypotheses, as in Ref. 41. It is not an optimal Neyman-Pearson test, however, because it optimizes over the unknown parameters of each hypothesis instead of integrating over known *a priori* probability density functions for them. This type of maximum-likelihood-based detection test is known in the literature as a Generalized Likelihood Ratio test [60]. The loss of detection power due to this sub-optimal implementation, however, is not expected to be very large because optimization of unknown parameters often tends to produce results similar to integration over their *a priori* probability densities.

The hypothesis test design is completed by selection of a detection threshold value γ_{th}. The γ statistic tends to be a large positive number when the signals are authentic because the optimal value of J_{nonsp} from Eq. (4b) tends to be small while the optimal value of J_{sp} from Eq. (5b) tends to be large due to the poor ability of the spoofed model in Eq. (2) to fit non-spoofed data as modeled by Eq. (1). If the signals are spoofed, then the situation reverses. The optimal value of J_{nonsp} tends to be large, that of J_{sp} tends to be small, and γ tends to be negative. Therefore, a good threshold value γ_{th} will be "near" 0 under some suitably defined definition of "near". If the calculated γ from Eq. (15) obeys $\gamma \geq \gamma_{th}$, then the signals are declared authentic. If $\gamma < \gamma_{th}$, on the other hand, then a spoofing attack is declared.

Standard hypothesis test design techniques pick the value of γ_{th} based on a desired upper bound on the probability of false alarm, i.e., the probability of achieving $\gamma < \gamma_{th}$ in the absence of spoofing. These techniques then calculate the probability of detection as the probability that $\gamma < \gamma_{th}$

under a spoofing attack scenario. Such calculations require *a priori* probabilities for the unknown quantities $\hat{\boldsymbol{r}}_{BA}$, β, β_{sp}, and $\Delta N_{BA}^{1},...,\Delta N_{BA}^{L}$ under the spoofed and non-spoofed assumptions.

Reference 41 avoided using *a priori* probability densities by working with worst-case probabilities of false alarm and worst-case probabilities of missed detection. It is not obvious how to modify the worst-case analyses of Ref. 41 for the present system due to the inclusion of the integer ambiguities in the problems in Eqs. (4a)-(4e) and (5a)-(5d).

The hypothesis detection threshold values γ_{th} used in the present paper have been designed by resorting to Monte-Carlo simulations. The inputs to these simulations are the quantities ρ_{BA}, λ, $\hat{\boldsymbol{r}}^{1}$, …, $\hat{\boldsymbol{r}}^{L}$, σ_{rcvr}^{1}, …, σ_{rcvr}^{L}, and σ_{mp} along with a random-number generator seed. Two simulations are conducted for a given analysis, a non-spoofed simulation and a spoofed simulation. Each non-spoofed simulation generates the set of non-spoofed single-differenced carrier phases $\Delta\phi_{BA}^{1}$, …, $\Delta\phi_{BA}^{L}$ for N_{MC} different cases by using the non-spoofed signal model in Eq. (1). The various unknown parameters in that equation are sampled from appropriate distributions using a random number generator. The antenna pair orientation vector $\hat{\boldsymbol{r}}_{BA}$ is sampled from a uniform distribution on the unit sphere. The unknown line bias β is sampled from a flat distribution over the range [-19797.31, +19797.31] radians. Each ΔN_{BA}^{j} is sampled from a flat integer distribution over the range [-5000, +5000]. Each random noise term n_{mpBA}^{j} and n_{rcvrBA}^{j} is sampled from a zero-mean Gaussian distribution with the respective standard deviations σ_{mp} and σ_{rcvr}^{j}. The resulting $\Delta\phi_{BA}^{1}$, …, $\Delta\phi_{BA}^{L}$ values have integer multiples of 2π subtracted from them so that their final values lie in the range $[-\pi,\pi]$. They are input to the spoofing detection calculations in Eqs. (4a)-(4e), (5a)-(5d), and (15) in order to produce N_{MC} samples of the γ detection statistic. The result is a histogram approximation of the probability distribution of γ under the non-spoofed hypothesis. In the limit of large N_{MC}, this histogram is an exact representation of the true non-spoofed distribution γ, $p(\gamma|H_0)$. Note that the random distributions used for β and for each ΔN_{BA}^{j} are probably much wider than necessary to generate a good simulation of the true γ statistics, but there is no expected drawback from using these widened distributions.

The second Monte-Carlo simulation for a given analysis yields an approximation of the spoofed-case probability density of γ. It is similar to the corresponding non-spoofed simulation. The only difference is that the spoofed signal model in Eq. (2) is used to generate the simulated single-differenced carrier phase measurement set $\Delta\phi_{BA}^{1}$, …, $\Delta\phi_{BA}^{L}$ for each of the N_{MC} simulation cases. It uses random number generators in order to sample the unknown spoofed line bias β_{sp} from the sum of a flat distribution over the range [-19797.31, +19797.31] radians and a zero-mean Gaussian distribution with standard deviation σ_{mp}. Each ΔN_{BA}^{j} is sampled from the flat integer distribution in the range [-5000, +5000], and each random noise term n_{rcvrBA}^{j} is sampled from a zero-mean Gaussian distribution with standard deviation σ_{rcvr}^{j}. All of the other operations are the same as for the non-spoofed case, and the result is a histogram of N_{MC} simulated samples of γ for the spoofed case. It is a good approximation of the γ probability density function under the spoofed hypothesis, $p(\gamma|H_1)$.

Fig. 4 shows the simulated non-spoofed and spoofed probability density histograms for a representative case. Important parameters for this case are the antenna separation ρ_{BA} = 0.14 m, the number of GPS satellites L = 7, their GDOP: 2.4, and the range of the received carrier-to-noise ratios for the 7 signals in the two antennas C/N_0: 34.1 to 49.7 dB-Hz. The number of simulated Monte-Carlo detection statistics for the two cases is N_{MC} = 10,000. The receiver thermal noise standard deviation is $\sigma_{rcvr}^{j} = [B_{PLL}/(C/N_0)^{j}]^{0.5}$, where B_{PLL} = 2.6 Hz is the Phase-Lock Loop bandwidth. The single-differenced multipath error standard deviation is σ_{mp} = 0.33 rad, which is equivalent to 0.01 m of single-differenced range error at the GPS L1 frequency.

The blue histogram in Fig. 4 is the γ probability density histogram without spoofing $p(\gamma|H_0)$, and the red curve is for the spoofed case $p(\gamma|H_1)$. The red curve is very narrow, so narrow that it appears to be a Dirac delta function in comparison to the wide blue distribution, and its peak extends above the top of the figure's vertical scale. The magenta vertical line is a candidate value for the spoofing detection threshold, $(\gamma_{th})_{candidate}$ = 1250. It clearly lies above all of the red spoofed-case histogram and below all of the blue non-spoofed-case histogram. Therefore, it seems likely that this threshold value achieves a probability of false alarm P_{FA} that is below $1/N_{MC}$ = 0.0001 and a probability of missed detection P_{MD} that is also below $1/N_{MC}$ = 0.0001. Thus, this is a powerful spoofing detection test.

It would be expensive to conduct a new Monte-Carlo analysis in order to design a good spoofing detection threshold γ_{th} for any possible spoofing scenario. It would be best to develop analytic expressions for the spoofed and non-spoofed probability density functions $p(\gamma|H_1)$ and $p(\gamma|H_0)$. $p(\gamma|H_0)$ could then be used to solve for γ_{th} in

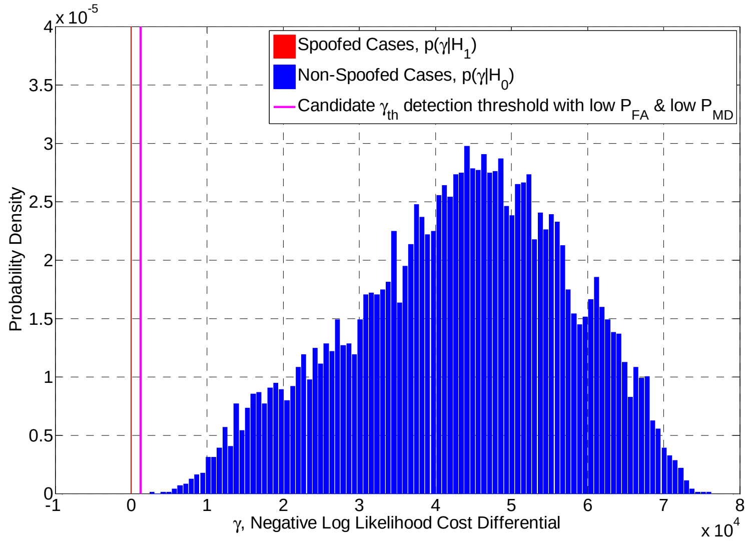

Fig. 4. Simulated probability density of spoofing detection statistic γ for spoofed (red) and non-spoofed (blue) hypotheses for a typical spoofing detection scenario.

order to achieve a given P_{FA}, and this γ_{th} value could be used with $p(\gamma|H_1)$ in order to compute the corresponding P_{MD} [41,59]. If it proves too difficult to derive appropriate analytic expressions for $p(\gamma|H_0)$ and $p(\gamma|H_1)$ or to calculate the needed P_{FA} and P_{MD} integrals, then it may be possible to perform an off-line Monte-Carlo analysis to develop a table of suitable γ_{th} detection thresholds. There might be only two principal inputs to this tabulated function: the satellites' weighted GDOP and their average C/N_0. This is an open issue that remains to be studied.

IV. A PLL TO ESTIMATE SINGLE-DIFFERENCED CARRIER PHASE FOR A SWITCHED-ANTENNA SYSTEM

A special PLL is needed in order to properly track the received carrier phase for the switched-antenna version of the two-antenna system, the one depicted on the left-hand side of Fig. 1. The antenna switching causes step changes in the received beat carrier phase, and a standard PLL would experience ringing at each step, as depicted in Fig. 5. This ringing is problematic for three reasons. First, it reduces tracking robustness. Second, it reduces the accuracy of the tracked beat carrier phase. Third, it does not allow for accurate determination of the phase jump caused by an antenna switch, which constitutes the $\Delta\phi_{BA}$ value that is needed for input to the spoofing detection calculation.

All three of these problems can be solved by developing an augmented PLL that is designed to deal with the phase jumps explicitly. The antenna switching times can be fed into an augmented PLL because they are available from the switch control line shown in the left-hand system of Fig. 1. The augmented PLL would store and update an estimate of the magnitude of this jump, $\Delta\phi_{BA}$. For all accumulation intervals corresponding to data from Antenna B, the augmented PLL discriminator would subtract the PLL's estimate of $\Delta\phi_{BA}$ from the standard atan2 discriminator before feeding it back through the PLL loop filter. These Antenna-B intervals correspond to the upper switched portions of the blue curve in Fig. 5. The PLL would update its estimate of the single-differenced phase $\Delta\phi_{BA}$ at the end of each switching period. Its update would be based on the residual ringing in the discriminator during the preceding switch period.

This augmented PLL would eventually settle to a correct estimate of $\Delta\phi_{BA}$ and zero ringing at the switch times, as depicted in Fig. 6 and Fig. 7. Note how the ringing at the antenna switching times in Fig. 6 decays with time and how the estimated single-difference carrier phase $\Delta\phi_{BA}$ approaches the true value in Fig. 7. The antenna switching frequency is a system design parameter, and the decay time constant of the $\Delta\phi_{BA}$ estimation error is a PLL design parameter, similar to its loop bandwidth. A low switching frequency and a slow decay of the $\Delta\phi_{BA}$ error have been used for the example in Fig. 6 and Fig. 7 purely for illustrative purposes. An actual spoofing

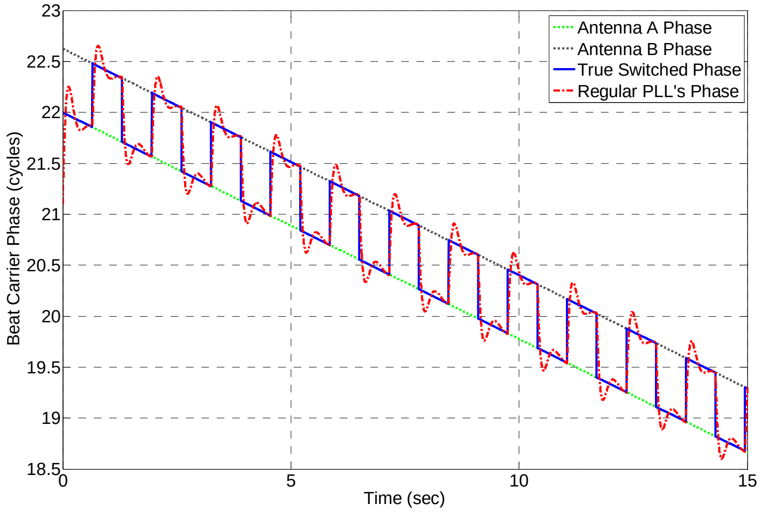

Fig. 5. Response of a standard PLL to the periodic phase jumps of a switched-antenna system.

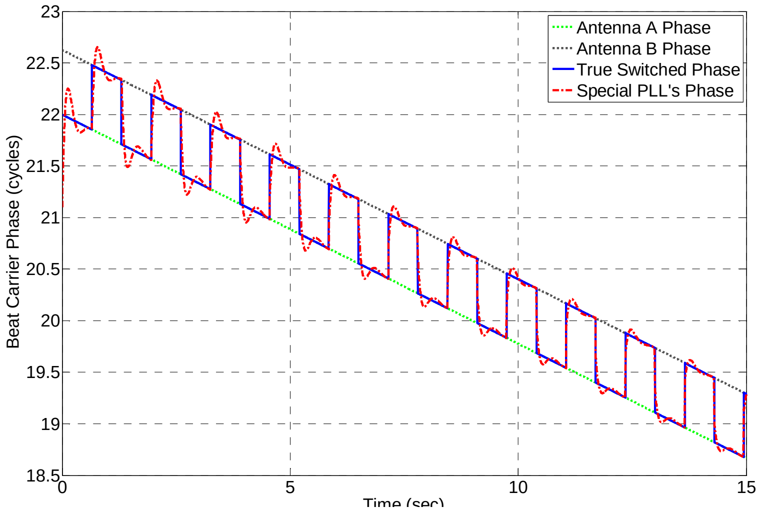

Fig. 6. Special PLL's response to switched-antenna phase time history with periodic jumps.

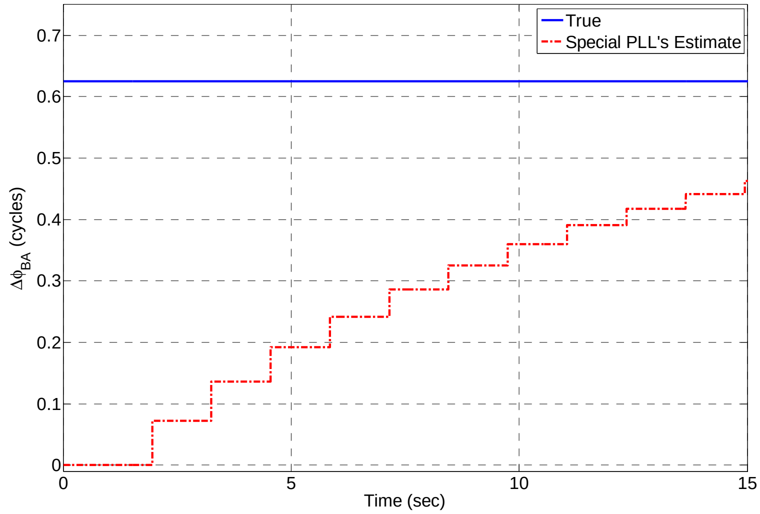

Fig. 7. True and estimated $\Delta\phi_{BA}$ time histories for special PLL operating on switched-antenna data.

detection system would likely use a higher switching frequency and a faster error decay time constant in order achieve fast detection of a spoofing attack.

A prototype of this special PLL has been developed, but more work remains to be done before it can be described in detail. It has not yet been implemented in a real-time system that also has an actual switched antenna and a means of supplying the switching times to the special PLL. A remaining hurdle in the development of this PLL is the need for techniques to deal with accumulation intervals that overlap switch times and with phase ambiguities caused by unknown navigation data bits. The architecture, analysis, and evaluation of this system are left as potential subjects for a future publication.

The switched-antenna concept could be generalized to employ more than 2 antennas. The specialized PLL would need a unique identifier for each reported switching time, one that indicated the two antennas involved in the given switch. Such a system could be used in conjunction with a multi-antenna spoofing detection methods described in Refs. 27, 33, and 40. A switched-antenna version of the multi-antenna system would have a single RF front end and a single receiver channel per satellite, which would reduce its hardware complexity. Like the system of Refs. 27, 33, and 40, its outputs would enable determination of the full unit direction vector to each received signal, with determination being made in a coordinate system defined relative to the multi-antenna system.

V. OFFLINE AND LIVE-SIGNAL TESTING OF THE TWO-ANTENNA SPOOFING DETECTION SYSTEM

Offline and live-signal testing has been performed using a prototype version of the two-antenna system. This section describes the system that has been tested, the testing equipment that has been used, and the experiments that have been performed. It then reports on the test results.

A. Blue Team: The Prototype Two-Antenna Spoofing Detection System used in Real-Time Tests

The prototype spoofing detection system used in this study is an example of the two-antenna configuration depicted on the right-hand side of Fig. 1. Its two antennas connect to two independent RF front ends that run off of the same reference oscillator. These RF front-ends provide input to two independent receivers that track each signal using a Delay-Lock Loop (DLL) and a PLL. The beat carrier phases of the PLLs, ϕ_A^1, …, ϕ_A^L from the Antenna-A receiver and ϕ_B^1, …, ϕ_B^L from the Antenna-B receiver, are output to the spoofing detection algorithm, which differences them in order to compute the inputs $\Delta\phi_{BA}^1$, …, $\Delta\phi_{BA}^L$ for the spoofing detection algorithm.

The elements of the prototype system are shown in Fig. 8 and Fig. 9. The two-antenna system consists of two GPS patch antennas mounted on a single ground plane with a spacing of 0.14 m. The two RF front ends are Ettus Research Universal Software Radio Peripherals (USRPs), and their common oscillator is an Ovenized Crystal Oscillator (OXO). The two receivers' digital signal processing functions are implemented in real-time software radio receivers (SWRX) that run in parallel on a single Linux laptop. They are written in the C++ programming language. The spoofing detection calculations are performed on the same Linux laptop using algorithms that have been encoded in MATLAB.

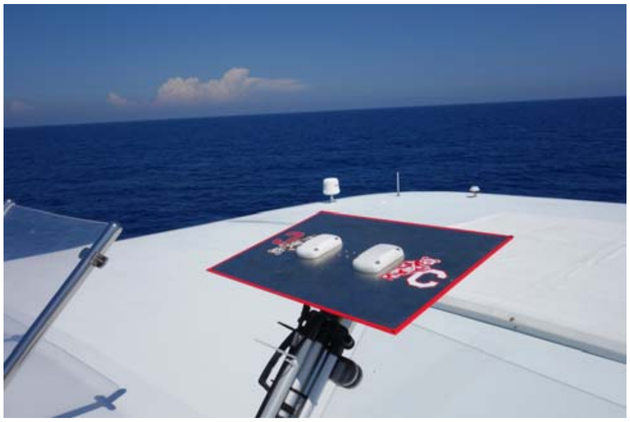

Fig. 8. The two antennas of the prototype spoofing detection system mounted on a common ground plane.

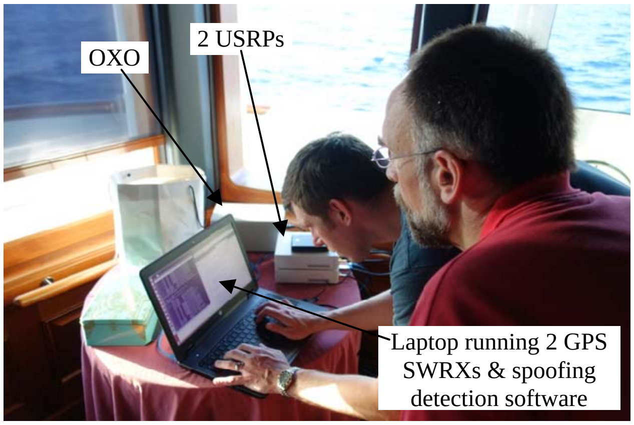

Fig. 9. Signal processing hardware of the prototype spoofing detection system.

One of the key features of this system's architecture is the ability of its real-time software radios' C++ code to call the spoofing detector's MATLAB tic function and to pass carrier-phase and other relevant data to the tic function. This capability shortened the implementation and test cycle for the prototype system by eliminating the need to translate the original MATLAB versions of the spoofing detection algorithms into C++. This code architecture enabled rapid re-tuning and re-design of the spoofing detection calculations. This capability was exploited during the course of live-signal testing.

The MATLAB spoofing detection software package produces a display for purposes of real-time signal authentication during spoofing detection tests. Several different versions of this display have been generated. Fig. 10 shows the version of the display that was used for this study's culminating live-signal tests. The 4 panels of this figure are, clockwise from top left, a) the spoofing detection statistic time history $\gamma(t)$, b) four diagnostic time histories that include time histories of the number of satellites used for spoofing detection $L(t)$ (blue asterisks), their corresponding GDOP(t) values (magenta o's), the time increment between spoofing detection tests $\Delta t_{spf}(t)$ (green dots), and the compass heading $\psi(t)$ as determined from the two-antenna $\hat{\boldsymbol{r}}_{BA}$ non-spoofed-case solution (black dots), c) Compass display, and d) time history of GPS PRN number availability. All of these displays are updated in real-time. The upper-left, upper-right, and lower-left time history plots scroll along their horizontal time axes in order to keep the most recent 4.5 minutes of data available. The lower-right compass updates each time a new spoofing detection calculation is performed. The green dots in the upper-left plot indicate that the time between spoofing detections, Δt_{spf}, is nominally 1 second, though sometimes the gap is longer due to lack of sufficient number of validated single-differenced carrier phases to carry out the calculation. Thus, the nominal update time for all of the plots in this display is 1 second. Faster updates are possible with the MATLAB software, but Δt_{spf} was deemed sufficiently fast for this study's experiments.

The most important panel in Fig. 10 is the upper-left spoofing detection statistic time history. The magenta plus signs on the plot show the spoofing detection threshold chosen for this case, γ_{th}. The computed γ values are plotted as green o's if they lie above γ_{th} and as red asterisks if they lie below. If the current γ value is above γ_{th}, then the message '`GPS Signals Authenticated`' is displayed on the plot. If the current γ value is below γ_{th}, then the displayed message switches to the spoofing alert: '`GPS SPOOFING ATTACK DETECTED!`', as in Fig. 10. Thus, the vertical level of the last plotted γ point, its plotted color and symbol type, and the presence of one or the other of the above two messages provide triple redundancy in alerting the operator to the authentication status of the GPS signals. A later version of the display, not shown, uses six panels. It includes all of the information in Fig. 10, and it has a separate panel devoted to a large circle that acts as a highly visible authentication status indicator: It is green if the signals are authenticated and red if a spoofing attack has been declared. The use of MATLAB to implement the real-time display enabled rapid reconfiguration to add such features.

The other 3 panels in Fig. 10 proved helpful in diagnosing this prototype system's performance. A low L value (near 4) or a high GDOP value in the upper-right panel indicated poorer reliability of the spoofing detection

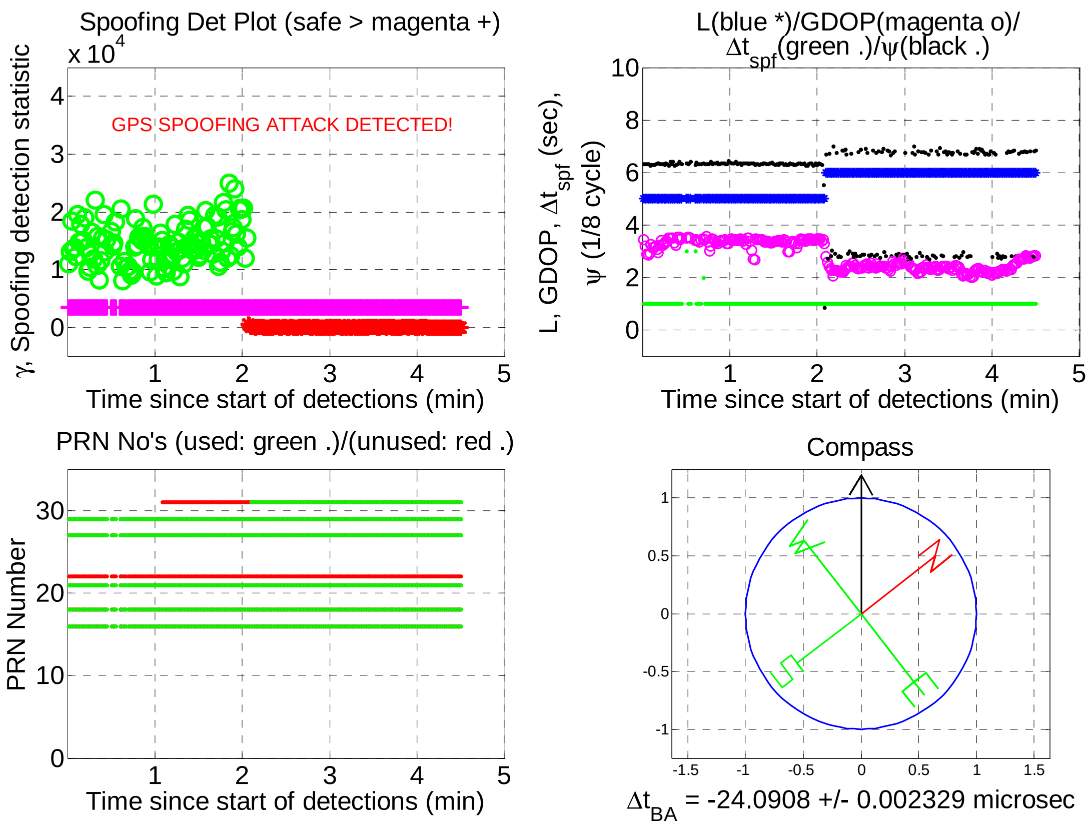

Fig. 10. Example snapshot of spoofing detection software's real-time display panel as used during live-signal attacks.

calculations. A correct compass heading in the absence of spoofing provided a check on the system. During spoofing attacks, the compass heading became jumpy, thereby providing another possible indicator of inauthentic signals. Note how the black dots of the $\psi(t)$ plot in the upper-right panel start jumping back and forth between two wrong compass headings just after $t = 2$ min, which is the same time that the $\gamma(t)$ plot in the upper-left panel dives below the γ_{th} spoofing detection threshold.

The vertical scale of the lower-left panel lists the possible GPS PRN numbers. The presence of a green or red dot at the level corresponding to a given PRN number indicates that one or both receivers is seeing something from that satellite at the corresponding time. If the dot is red, then the returned data are incomplete or are deemed to be insufficiently validated for use in the spoofing detection calculation. If the dot is green, then the data from that PRN have been used in the detection that has been carried out at that time. A close comparison between the upper-right $L(t)$ plot (blue asterisks) and the lower-left panel demonstrates that L is equal to the number of green dots at any given time. Note how $L(t)$ jumps from 5 to 6 in the upper-right panel at the same time that the dots for PRN 31 switch from red to green in the lower-left panel, just after $t = 2$ min. A red dot in the lower-left panel can indicate many things. Perhaps only one of the two receivers is tracking that signal. Alternatively, its C/N_0 value in one or both antennas may be too low to consider the corresponding phase data to be reliable.

The time difference Δt_{BA} shown below the lower-right panel of Fig. 10 is the estimated timing error between the two USRPs' sample clocks. This has been computed based on differences of the DLL PRN code start times between the two receivers for all tracked signals. This difference should be on the order of 1 nsec for properly functioning USRPs and their associated drivers. Unfortunately, it is likely that initial packet losses in the USRP-to-computer data transmission gave rise to initial errors. The use of MATLAB to implement the spoofing tic function enabled quick implementation of a real-time estimation algorithm for Δt_{BA}, which allowed compensation for it in the spoofing detection calculations. The spoofing detector produced nonsensical results prior to implementation of this compensation. A switched-antenna system, as shown on the left-hand side of Fig. 1, would not have such timing issues.

Another feature of the prototype spoofing detection system is its ability to record the wide-band RF data from its two antennas. For each spoofing scenario, the raw samples from both USRPs were recorded while the real-time software receiver was performing its signal processing operations and while the real-time spoofing detector was doing its calculations. These recorded data streams will allow off-line analysis and off-line testing of a re-tuned or completely re-designed spoofing detection system.

B. Red Team: The Receiver/Spoofer used in the Real-Time Tests

The spoofing detection system described in this paper has been tested against the latest version of receiver/spoofer of Refs. 2, 3, 5, and 28. Its attack strategy is to overlay the spoofed signal on top of the true signals, ramp up the power to capture the receiver tracking loops, and finally drag the pseudorange, beat carrier phase, and carrier Doppler shift off from their true values to spoofed values. The pseudorange part of a spoofing attack is depicted in Fig. 11, which shows the cross correlation of the receiver's PRN code replica with the total received signal (blue solid curve), the receiver's early, prompt, and late correlations (red dots), and the spoofer signal (black dash-dotted curve). In the top plot, the spoofer has zero power, and the receiver sees only the true signal. The second and third plots show the spoofer ramping up its power while maintaining its false signal in alignment with the true signal. It is able to achieve this alignment by exploiting its own reception of the true signal and its knowledge of its antenna geometry relative to the victim receiver's antenna. The spoofer power in the middle/third plot is sufficient to capture control of the 3 red dots of the receiver's DLL. In the fourth and fifth plots, the spoofer initiates and continues a pseudorange drag-off, which is an intentional falsification of the pseudorange as measured by the victim receiver's DLL.

The spoofer performs drag-off simultaneously on all spoofed channels in a vector spoofing attack that maintains consistency of all spoofed pseudoranges. After the initiation of drag-off, the victim receiver computes a wrong position, a wrong true time, or both, but the residual pseudorange errors in its navigation solution remain small. Therefore, this type of attack is not detectable by traditional pseudorange-based RAIM calculations [1].

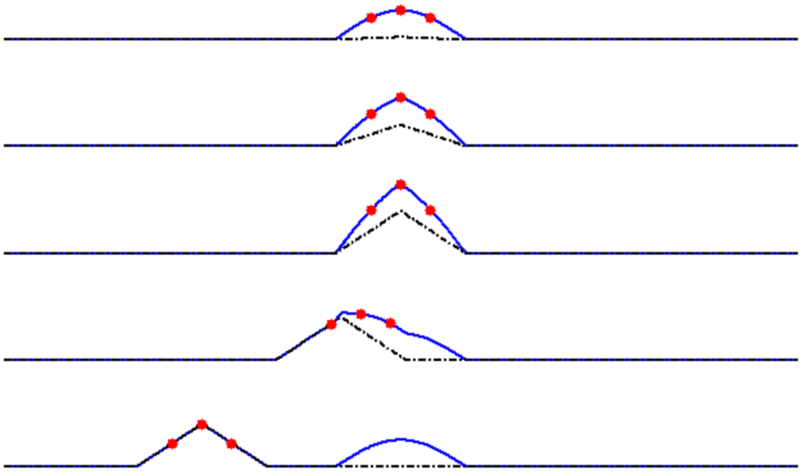

Fig. 11. Spoofer/receiver attack sequence as viewed from a channel's code offset cross-correlation function. Spoofer signal: black dash-dotted curve; sum of spoofer and true signals: blue solid curve; receiver early, prompt, and late correlation points: red dots.

The receiver spoofer hardware consists of a GNSS reception antenna, the receiver spoofer signal processing unit, and the spoofer transmission antenna. The reception and transmission antennas must be sufficiently separated, both spatially and in terms of their orientations and gain patterns, so that the receiver/spoofer does not self spoof. Fig. 12 shows the hardware set-up of the receiver/spoofer system that was used in this study's live-signal attacks. The system was mounted on a superyacht named the White Rose of Drachs (WRoD), depicted in Fig. 13. The upper-left panel of Fig. 12 shows the GPS reception antenna of the receiver/spoofer. It was positioned on the rear upper deck of the WRoD. The upper-right panel shows the receiver/spoofer transmission antenna mounted on the forward part of the WRoD sun deck. It was a directional antenna, and it was pointed forward at the WRoD GPS antenna. The spoofing detector antenna pair was positioned near the defended WRoD antenna, with both being located on the roof of the WRoD bridge. The orientation of the spoofing transmission antenna, combined with its remote location from the receiver/spoofer's reception antenna, ensured that the spoofer did not self spoof. The spoofer electronics, shown in the lower panel of Fig. 12, were located amidships on the upper deck of the WRoD.

The receiver/spoofer requires tuning of its transmission power levels. If the power is too high, then its spoofing attacks will be too obvious. Also, a very high transmitted power could saturate the front-end electronics of the intended victim, which would cause it to jam the system rather than spoof it. If the transmitted power is too low, then it will not capture the victim's tracking loops, and its spoofing attack will fail. The proper power level depends on the gain patterns of the spoofer transmission antenna and the victim receiver antenna and on their relative geometry.

The spoofing detection tests described here were conducted after initial tuning adjustments of the spoofer transmission power. Tuning was carried out through the digital selection of spoofer power levels and through proper choice of analog attenuators in the signal path from the spoofer electronics to its transmission antenna. The power level was tuned based on indicators of spoofer capture

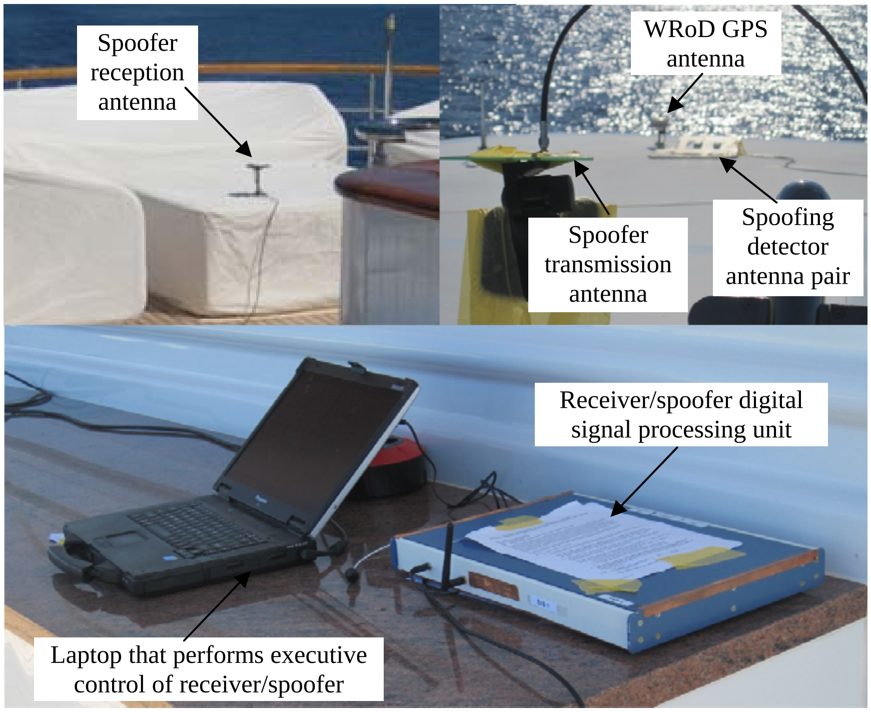
Spoofer reception antenna
WRoD GPS antenna
Spoofer transmission antenna
Spoofing detector antenna pair
Receiver/spoofer digital signal processing unit
Laptop that performs executive control of receiver/spoofer

Fig. 12. The components of the receiver/spoofer hardware as used for live-signal spoofing detection experiments.

Fig. 13. The White Rose of Drachs, the superyacht on which this study's live-signal spoofing detection tests were conducted.

of the WRoD GPS receiver.

C. Spoofing Attack Test Scenarios

Three sets of tests have been conducted in order to develop and evaluate this paper's spoofing detection system. The first tests started by recording wideband RF GPS L1 data using USRPs. These data were post-processed in two software receivers that recorded the outputs of their signal tracking loops. Afterwards, the MATLAB spoofing detection calculations were run using the recorded tracking loop data as inputs. The data recording and software receiver signal processing were both carried out at UT/Austin in March 2014. The subsequent spoofing detection calculations were carried out at Cornell, also in March 2014. Non-spoofed scenarios were recorded by placing the two antennas outdoors. Spoofed cases were simulated by placing the two antennas in an RF enclosure along with a GPS antenna that re-radiated the L1 band as received from a roof-mounted antenna. These tests proved the efficacy of the spoofing detection algorithms. They did not, however, test the performance of the system during the transition from non-spoofed to spoofed signals that takes place at the initiation of a spoofing attack.

The second set of tests was carried out using the first real-time version of the system. Before they could be carried out, the MATLAB spoofing detection calculations had to be re-packaged into a tic function that could be called in real-time by the C++ real-time software receivers that process the USRP data from the two antennas. The initial tests of this system were conducted at UT Austin using real-time replay of the recorded USRP data from the first tests. These tests were used to refine the design of the MATLAB code and of its interface to the C++ SWRXs so that real-time operation could be achieved. Successful real-time tests were performed using a 3.33 Hz frequency for the spoofing detection calculations of Eqs. (4a)-(4e), (5a)-(5d), and (15).

The first live tests of the real-time system were conducted at Cornell University in Ithaca, NY in mid June 2014. The non-spoofed case was tested by operating the system on the roof of Cornell's Rhodes Hall. The spoofed case was also tested on the roof of Rhodes Hall, except that the output of a single antenna was sent through a signal splitter and used to drive both USRPs. According to theory, this set-up should produce the same response as would be produced by a spoofing attack. These rooftop tests were used to debug the USRP timing error that has been mentioned in conjunction with the lower-right panel of Fig. 10. The compass display in that same panel of Fig. 10 served to verify proper operation for non-spoofed cases on the

Rhodes Hall roof, and the spoofing detection statistic in the upper-left panel confirmed proper operation for both the spoofed and non-spoofed cases. This set of tests also was unable to probe the system's performance at the onset of a spoofing attack, before the signal drag-off depicted in the lower two traces of Fig. 11.

The final set of tests was conducted aboard the WRoD. The spoofing detection experiment team boarded the WRoD on June 24 and 25, 2014 in Cap d'Ail, France, which is the next town west of Monaco on the Riviera. While in port, the team re-conducted the same set of tests that had been conducted on the Rhodes Hall rooftop in Ithaca, NY. These tests confirmed that the spoofing detection equipment had survived its travels. Live-signal attacks could not be conducted in port because live-signal spoofer transmissions in the GPS L1 band are allowed for scientific testing purposes only if conducted in international waters.

Late on June 26, 2014, the WRoD set sail for a cruise around Italy to Venice, where it arrived on the afternoon of June 30th. Three separate sets of live-signal spoofing attacks and detection tests were conducted on June 27th, 28th, and 29th. The initial attacks and detection tests on June 27th were conducted mostly for purposes of selecting antenna locations and tuning the spoofing power of the receiver/spoofer. Later attacks explored other aspects of spoofing and detection.

The power tests on June 27th needed a means to decide whether a given attack had captured the tracking loops of the WRoD GPS receiver, which was a Litton LMX 420. The strategy for confirming capture was to perform a noticeable drag-off after the initial attack. The drag-off needed to be in an obvious direction so that it would be clear whether the Litton receiver was tracking the true WRoD course or the spoofed course. After some initial testing, it was determined that a vertical drag-off would provide the most obvious indication of a successful capture. Successful attacks dragged the Litton LMX 420's reported altitude as high as 5000 m. In practice, a reported altitude above 25 m was usually sufficient to confirm the success of a given spoofing attack.

The tests that evaluated spoofer and spoofing detector antenna placements relative to the WRoD GPS antenna were also important to achieving sensible results. Various placements were tried with varying degrees of success in capturing the WRoD receiver and in detecting the capture using this paper's new system. The most successful relative geometry for the spoofer antenna, the spoofing detector antennas, and the WRoD antenna was the configuration depicted in the upper-right panel of Fig. 12.

Note that the question of placement of the spoofing detector antennas relative to the defended WRoD antenna is atypical of likely real-world detection scenarios. It is expected that a real-world spoofing detector will be integral with the defended GNSS receiver. Thus, there will be no question about the relative placement of their antennas.

The culminating live-signal attack on June 29th involved a 50 minute spoofing scenario in which the attacker took the WRoD from the Adriatic to the coast off of Libya. The scenario's long distance and short duration required a mid-course speed in excess of 900 kts. This spoofing scenario was designed in the simplest possible way, by taking a straight-line course in WGS-84 Cartesian coordinates from the true WRoD location to the spoofed location off of Libya. This course took the spoofed yacht position across the Italian and Sicilian land masses and below the Earth's surface to a maximum depth of more than 23 km.

Obviously, the WRoD was physically unable to execute this maneuver. Its crew would not have needed this paper's spoofing detection system in order to realize that its GPS receiver was returning false readings. The main points of this last test were to dramatize the potential errors that can be caused by a spoofer and to check whether the spoofing detector could continue to function under these drastic conditions.

Fig. 14 highlights this unusual scenario by presenting views of two displays on the WRoD bridge, as photographed during the attack. The GPS display in the upper panel shows the speed 621 kts and the altitude -7376 m. The chart display in the lower panel shows the yacht on (or rather, below) dry land and halfway across the "insole" of Italy's boot. It also shows a tremendously long velocity vector, one that extends beyond the lower edge of the chart.

D. Spoofing Detection Test Results

As an example, consider some results from the dramatic Libya spoofing attack scenario. Fig. 15 plots various signal output time histories from the spoofing detection system. It illustrates the attack sequence and suggests means by which the spoofing detection system can be evaluated. Its upper panel plots the fractional portions of the two-antenna spoofing detector's single-differenced beat carrier phase time histories $\Delta\phi_{BA}^{1}$, …, $\Delta\phi_{BA}^{L}$ for the L = 7 tracked PRN numbers 16, 18, 21, 22, 27, 29, and 31. The figure's middle panel plots the amplitude time history of the 100 Hz prompt [I;Q] accumulation vector for PRN 16, as received at Antenna A of the detection system. The figure's bottom panel plots the PRN 16 carrier Doppler shift time history, as determined by the detection system's PLL for Antenna A.

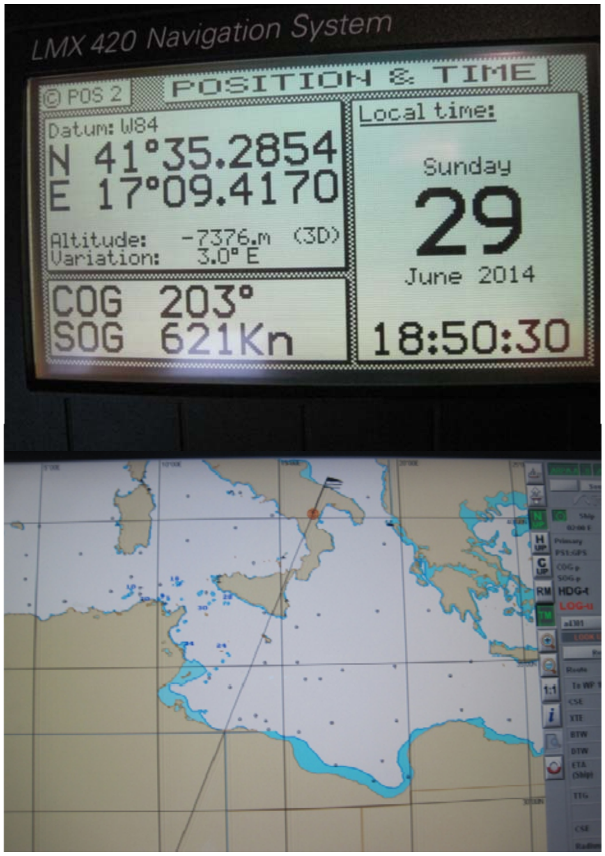

Fig. 14. The WRoD bridge GPS receiver display (top picture) and its GPS-driven chart (bottom picture) during the Libya spoofing scenario.

This was a strong attack in which the spoofer power was 10.7 dB higher than the power of the real signal for PRN 16. The other spoofed signals had power advantages over their corresponding true signals that ranged from 3.3 dB to 13.6 dB, and the spoofer's mean power advantage was 10.4 dB. Therefore, the onset of the spoofing attack at t = 196.1 sec is clearly indicated by the sudden jump in $(I^2+Q^2)^{0.5}$ on the middle panel. The upper panel shows a corresponding sudden coalescing of the $\Delta\phi_{BA}^{j}$ single-differenced beat carrier phases, which implies that this paper's spoofing detection algorithm should have been able to detect this attack.

The spoofer drag-off started at t = 321.5 sec, as evidenced by the sudden change in the slope of the carrier Doppler shift time history on the lower panel of Fig. 15. The period after the initial attack and before the drag-off is delimited by the vertical magenta and cyan dash-dotted lines. During this interval the spoofer was waiting to capture the receiver's tracking loops.

Note how the single-differenced $\Delta\phi_{BA}^{j}$ time histories in the upper plot of Fig. 15 appear to be somewhat more noisy during the interim pre-drag-off period of the attack than after the start of the drag-off at t = 321.5 sec. The grey dotted curve for PRN 27 is an exception because it becomes noisy again starting at about t = 450 sec due to decreased signal power. The increased noisiness of the PRNs' $\Delta\phi_{BA}^{j}$ time histories during the interim period is probably the result of interference between the true and spoofed signals, which are likely beating slowly against each other. The response of the spoofing detection algorithm during this phase is uncertain because this multipath-like beating between the two signals is not modeled by Eqs. (1) or (2). The detection system's response during such pre-drag-off attack periods was not tested until the live-signal attacks that were conducted aboard the WRoD.

Fig. 16 demonstrates the performance of this paper's spoofing detection algorithm for the Libya attack scenario. The upper panel of the figures is a repeat of the upper panel of the single-differenced beat carrier phase time histories from Fig. 15, except that they are plotted for a longer duration. Recall that these are the principal data that are input to the spoofing detection calculations in Eqs. (4a)-(4e), (5a)-(5d), and (15). The lower panel of Fig. 16 shows the $\gamma(t)$ spoofing detection statistic time history. It plots the same information that appeared in the upper-left panel of Fig. 10 during the corresponding real-time detection tests. At t = 196 sec $\gamma(t)$ is clearly above the blue dash-dotted spoofing detection threshold γ_{th}. At t

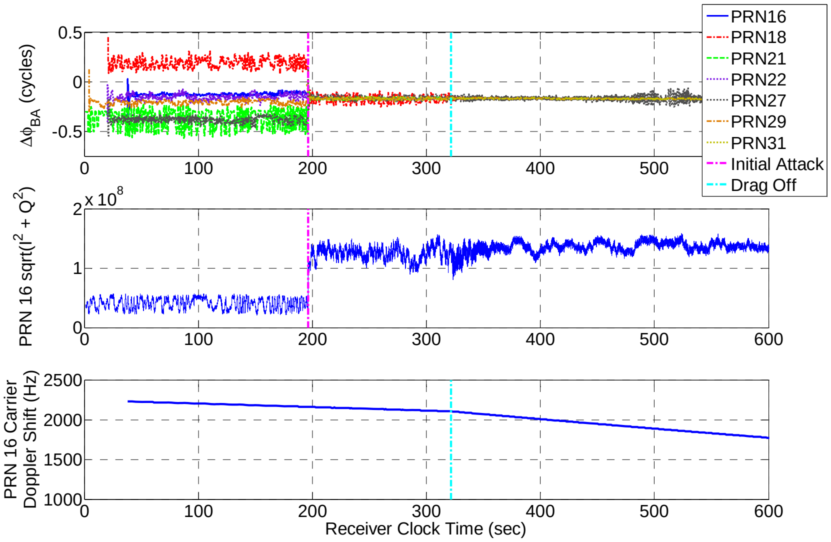

Fig. 15. Indicators of initial capture and drag-off during Libya spoofing attack, as measured by the spoofing detection receiver.

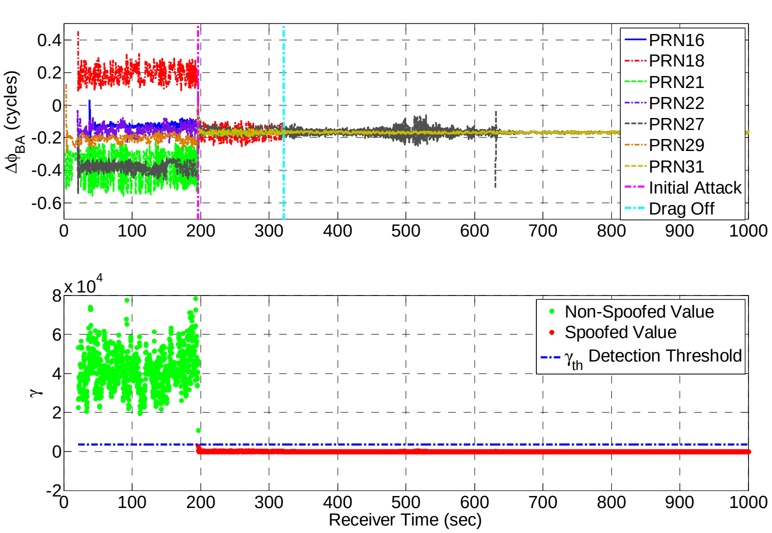

Fig. 16. Single-differenced carrier phase time histories and corresponding spoofing detection statistic time history for Libya spoofing attack scenario.

= 196.4 sec it is clearly below γ_{th}, which indicates a spoofing detection. It remains below γ_{th} for the duration of the attack. In this re-processed version of the detection calculations, $\gamma(t)$ has been updated at 5 Hz. Therefore, the earliest possible detection point would have been t = 196.2 sec, which is 0.1 sec after the onset of the attack. This point corresponds to the green dot in the lower panel of Fig. 16 that lies slightly above the blue dash-dotted γ_{th} line. Theoretically, the system might have detected the attack at this time, but the finite bandwidth of the two receiver's PLLs caused lags in the transitions of the single-differenced phases in the top plot, which led to the 0.3 sec lag in the detection of the attack. It is encouraging, however, that the spoofing detector worked well during the initial pre-drag-off phase of the attack, from t = 196.1 sec to t = 321.5 sec, despite the added noisiness of the single-differenced carrier phases in the top plot, which is likely caused by beating between the true and spoofed signals.

Fig. 17 plots the same quantities as in Fig. 16, but for a different spoofing attack. This one took place at about 13:30 UTC on June 27, 2014. It was a little less overt than the Libya attack. The power advantage of the spoofer ranged from 3.0 to 14.0 dB for the different channels with a mean power advantage = 9.2 dB. It was detected by this paper's system, as evidenced by the convergence of the single-differenced carrier phases at the onset of the attack at t = 397.5 sec. The spoofing detection statistic in the bottom panel of Fig. 17 dives near to the γ_{th} detection threshold at the onset of the attack and sometimes passes below it, but it does not stay permanently below the threshold until after the time of drag-off, i.e., after t = 531 sec. The large oscillations of the single-differenced carrier phases during the pre-drag-off initial capture interval from 397.5 to 531 seconds is likely due to beating between the true and spoofed signals. The largest variations occur for PRNs 12 and 31, which are the ones with the lowest spoofer power advantages, 3.2 and 3.0 dB, respectively. Apparently these oscillations cause $\gamma(t)$ sometimes to take on values slightly above γ_{th} during the interval 397.5 sec < t < 531 sec. Thus, the spoofing detector can experience problems in the initial phases of an attack; Eq. (2) does not exactly model the single-differenced phases at such times.

This spoofing detection case corresponds to the same problem parameters as have been used to generate the hypothesis test probability density histograms in Fig. 4. Thus, the horizontal axis of Fig. 4 corresponds to the vertical axis of the lower panel of Fig. 17. During the non-spoofed portion of Fig. 17, t < 397.5 sec, the green $\gamma(t)$ points in the lower panel of Fig. 17 range between 3×10^4 and 9×10^4, which is roughly consistent with the horizontal range of the blue histogram in Fig. 4. For the post-drag-off spoofed portion, t > 550 sec, the red $\gamma(t)$ points of Fig. 17 lie between -15 and 300. This is somewhat consistent with the horizontal range of the red histogram of Fig. 4, which lies between -20.5 and -7.5, but the intermittent presence of $\gamma(t)$ points as high as 300 in the bottom panel of Fig. 17 suggests that the

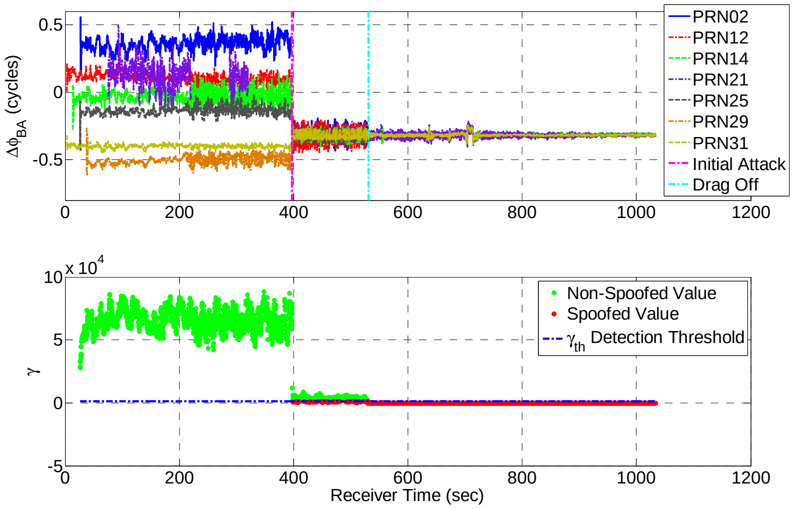

Fig. 17. Single-differenced carrier phase time histories (top plot) and spoofing detection statistic time history (bottom plot) for a spoofing attack with a slightly lower power advantage than the Libya attack.

spoofed-case signal model in Eq. (2) has some unresolved fidelity issues. Nevertheless, the spoofing detection threshold γ_{th} = 1250 still correctly identifies the spoofed cases after the drag-off. Given that spoofing prior to drag-off is benign in terms of whether the receiver determines a false position or time, the need to wait for drag-off for a definitive detection does not seem to constitute a serious weakness.

Note that the spoofer failed to capture the tracking loops of the WRoD's GPS receiver. This is surprising given the average spoofer power advantage of 9.2 dB above the true signals. It is conjectured that the WRoD GPS antenna had lower gain in the low-elevation direction toward the spoofer transmission antenna than did the spoofing detection system's antennas. A lower gain would reduce the spoofer power advantage in the WRoD receiver and could explain why the spoofer failed to deceive it.

Fig. 18 plots the single-differenced carrier phase time histories and the spoofing detection statistic for yet another attempted spoofing attack. This is the case of a failed attack. The upper plot of the single-differenced carrier phases fails to show convergence between the different channels at the onset of the attack because the spoofer was under-powered, leaving it no power advantage over the true signals. Its only effect on the single-differenced carrier phases was to make them more noisy due to beating of the weak spoofer signals against the stronger true signals, as during the time span 576 sec $< t <$ 712 sec.

Note that several short spoofing attacks were attempted during this test, with a succession of increasing power levels, but none with sufficient power to capture even the spoofing detector's tracking loops. The spoofer's most troublesome impact was that it temporarily confused the receiver's navigation data bit decoding on two channels. This confusion caused the single-differenced carrier phases for PRNs 06 and 24 to be flagged as being too unreliable to plot for the intervals 744 sec $< t <$ 838 sec for PRN 06 and 628 sec $< t <$ 935 sec for PRN 24. The spoofing detection algorithm reported an attack at about this same time. Note the low values of the red $\gamma(t)$ points on the lower panel of Fig. 18 during the interval 720 sec $< t <$ 850. It is unclear why the algorithm declared an attack when the L = 5 to 6 available single-differenced carrier phases did not converge during the alleged attack interval. This anomaly warrants further investigation. Although not a false alarm in the strictest sense of the word because spoofing was in progress, the spoofing criterion in Eqs. (4a)-(4e), (5a)-(5d), and (15) seems not to have been met. Therefore, an alert should not have been issued.

As for the previous case, the WRoD GPS receiver was not captured by any of these attacks. These attacks occurred at about 12:30 UTC on June 27, 2014, during the initial tuning of the spoofer power level.

Many additional spoofing attacks were carried out aboard the WRoD during the period June 27-29, 2014. The spoofing detector proved finicky. It took quite some time to get the spoofing detection two-antenna system positioned in a sensible place relative to the WRoD GPS antenna so as to be sensitive to nearly the same spoofing signals. In addition, the spoofing detector's GPS receiver tended to lose lock at the initiation of an attack, prior to signal drag-off. This was likely caused by the large power swings of the received signals due to beating of the true signals against the spoofed signals. This problem went away at higher spoofer power levels. When lock was lost, the software receiver would attempt to re-acquire the signal. Often a re-acquisition would succeed only after signal drag-off by the spoofer. Typically, the spoofing detector immediately detected the attack once it had re-acquired the spoofed signals that were no longer beating against the true signals due to having been dragged sufficiently far away from them, as in Fig. 11. Re-analysis of the recorded data indicated that poor PLL tuning may have caused the losses of lock during the initial attacks. Spoofing detection calculations carried out on the re-processed data have proved more reliable when implemented with a better PLL tuning.

Two attacks were carried out with only a subset of the

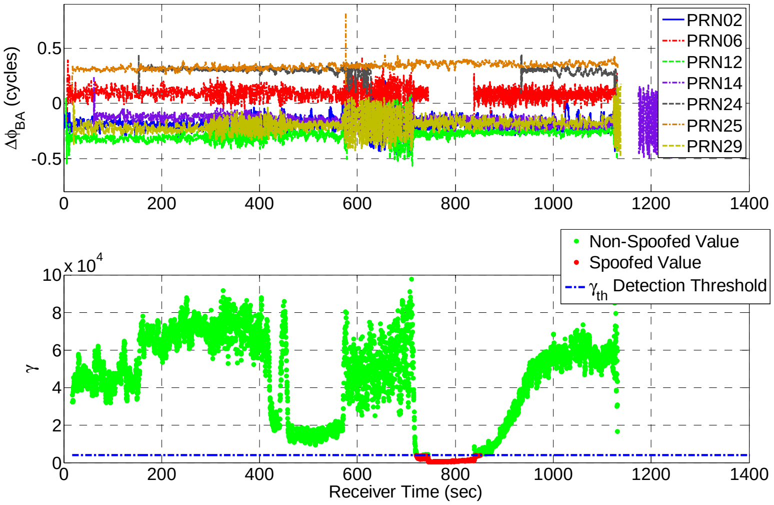

Fig. 18. Single-differenced carrier phase time histories (top plot) and spoofing detection statistic time history (bottom plot) for a failed spoofing attack with little or no power advantage relative to the true signals.

visible GPS satellites being spoofed. These tests were conducted in June 28, 2014 from UTC 17:25 to 18:10. The first test involved spoofing 7 of 9 visible satellites, and the second test spoofed only 4 of 9. The spoofing detection system had trouble maintaining signal lock during the initial part of the first attack. It subsequently re-acquired signals and was able to detect the attack successfully after re-acquisition. The first attack also succeeded in capturing the WRoD receiver's tracking loops as evidenced by spoofing of the yacht to climb off the sea surface. The second attack, the one with only 4 spoofed satellites, was not detected by this paper's prototype system, but it succeeded in deceiving the WRoD GPS receiver about its altitude. This latter result indicates a need to modify the detection calculations in Eqs. (4a)-(4e), (5a)-(5d), and (15) in order to allow for the possibility of partial spoofing. In their current form, they assume that all signals are either spoofed or authentic. Of course, in the partial spoofing case it may also be possible to use traditional pseudorange-based RAIM techniques to detect an attack, as per Ref. 1.

VI. POSSIBLE DIRECTIONS FOR FUTURE WORK

The data and results from this study suggest several possible avenues for fruitful development of improved GNSS spoofing detection and recovery systems. They are discussed below.

A. Improved Detection During Pre-Drag-Off Initial Phase of Spoofing Attack

The most obvious need for improvement of the current prototype system is its response to the initial phases of a spoofing attack. This is particularly important for the more subtle attacks in which the spoofer power advantage is not very large in comparison to the true signals. Two problems need to be solved in order to achieve reliable performance during this attack phase.

First, the tracking loops must be made more robust in the face of the large received power swings that are caused by the beating of the true signal against the spoofed signal. Simple re-tuning of PLL bandwidth might help. Alternatively, one might implement an FLL, yet keep track of the beat carrier phase in software. It would probably be helpful to allow the tracking loops for the two antennas to aid each other, especially for purposes of decoding unknown navigation data bits. Perhaps it would be possible to develop an advanced PLL that could explicitly recognize the existence of two signals that are beating against each other. Such a PLL might also be useful for tracking during extreme multipath or during equatorial ionospheric scintillation.

The second problem of spoofing detection during the pre-drag-off initial phase is that of how to calculate a good detection statistic. The beating of the true and spoofed signals makes the spoofed single-differenced phase model in Eq. (2) incorrect. A better signal model needs to be developed for this phase of an attack in cases where the spoofed signal's power is not much greater than that of the true signal.

B. Detection When only a Subset of the Signals are Being Spoofed

If only a subset of signals are being spoofed, then the algorithm developed in this paper may or may not work. The partial-spoofing cases reported in the previous section yielded successful detection when nearly all of the signals were spoofed. Detection was unsuccessful when less than half were spoofed. It might be possible to develop a combinatorial set of tests that would try different hypotheses about which signals were spoofed and which were not spoofed. An additional part of the test might involve a check of pseudorange consistency. Perhaps a test involving a combination of the single-differenced beat carrier phase observables and the pseudorange observables would offer an improvement to the present approach and to the pseudorange-only detection method that is defined in Ref. 1.

C. Advanced RAIM Techniques

Tracking-loop-level RAIM spoofing defenses might be useful during the initial pre-drag-off phase of a spoofing attack. Multipath estimation techniques like those of Ref. 61 might be helpful in characterizing the two received signals, the true signal and the spoofed signal. This might be done by considering both code-phase differences, which appear as distortions along the correlation early/late offset axis, and carrier-phase differences as seen at multiple antenna locations. Various criteria could be applied to the resulting estimates of the two distinct signals in order to determine whether the second signal is truly passive multipath or an active signal from a malicious spoofer. Earlier-signal/later-signal power comparisons might help, and the two components' single-differenced carrier phases between the two antennas might help too, should they prove to be uniquely determinable.

Another simpler RAIM-type technique is that of RF power monitoring, perhaps using the automatic gain control signal of the receiver's RF front-end, as suggested in Refs 32 and 45. Fig. 19 suggests that a simple power-based defense would have worked for some of the WRoD spoofing detections. It shows three histograms of raw RF front-end samples at the output of the Antenna-A USRP. At the top of each panel the root-mean-square value of the samples is reported as the histogram's σ value. These σ values are clearly different for the three cases. The non-spoofed case, the one in the left-hand panel, has $\sigma =$ 7100, and this is clearly the narrowest histogram. The weak spoofing that occurred in the middle panel widened this histogram, increasing its σ to 8200. This increased σ

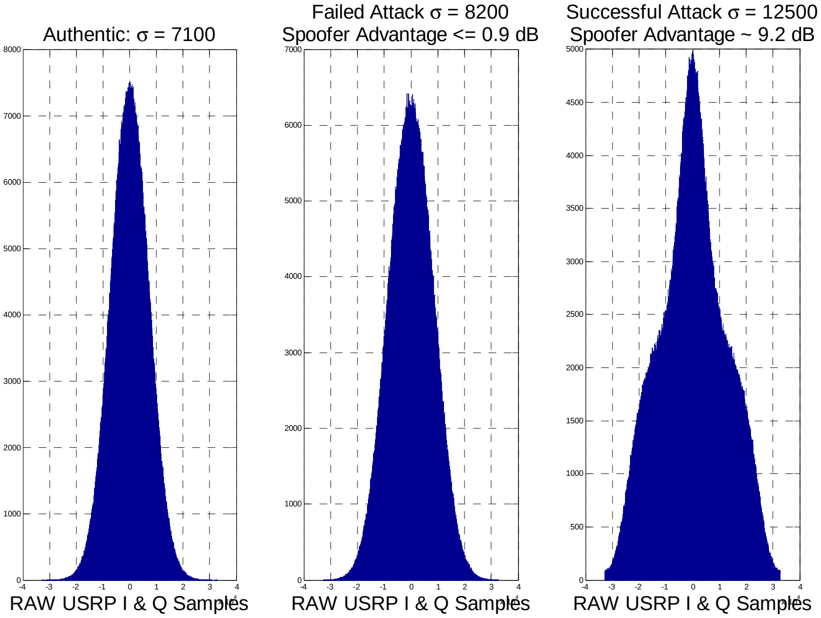

Fig. 19. Histograms of USRP raw RF samples for three different scenarios: no-spoofing (left-hand panel), low-power spoofing that failed to capture tracking loops (middle panel), and high-power spoofer that succeeded in capturing tracking loops (right-hand panel).

level might have been used to detect the spoofing attack, despite the fact that this attack is the unsuccessful one that corresponds to Fig. 18. Only the right-hand panel of Fig. 19 corresponds to a successful spoofing attack, the one associated with Fig. 17. The presence of the spoofing signals is apparent from its widened and distorted histogram and from its much larger σ value of 12500. Note, however, that jamming or solar radio bursts have the potential to trigger a spoofing false alarm for such a detection system. Also, it is likely that more subtle spoofing attacks could be mounted, ones that produced lower σ values and that captured all of the victim receiver's tracking loops. The attacks used in the present study were not the most sophisticated ones that are possible with the receiver/spoofer that was employed.

D. A Real-Time Prototype of the Switched-Antenna Version of this System

The switched-antenna system shown in the left-side panel of Fig. 1 could be developed into a working real-time prototype. It should yield a simpler, lower-power system. Its use of a single tracking loop for the two antennas might eliminate some of the uncertainty associated with the single-differenced carrier phase during the initial pre-drag-off portion of an attack.

E. Detection of a Spoofer that uses Multiple Transmission Antennas

This paper's detection algorithm is strongly dependent on the assumption that the spoofer transmits all of its false signals from a single antenna, as shown in Fig. 3. An advanced spoofer might transmit from multiple antennas. A detection test for such a spoofer might best be posed as an M-ary hypothesis test, with H_0 = no spoofing, H_1 = spoofing from 1 transmitter, H_2 = spoofing from 2 transmitters, etc. In the limit of a spoofer which uses L transmitters that have the correct geometry relative to the victim, it will be impossible to detect the attack by any generalization of the present techniques. If the number of spoofing transmitters is significantly lower than L, however, then attack detection may be straightforward.

F. Re-Acquisition of True Signals to Recover from a Spoofing Attack

The system presented in this paper only detects spoofing. Once an attack has been declared, a GNSS receiver is left only with the information that its position and time solutions are not authentic. Preliminary analysis of the recorded wideband WRoD spoofing data has demonstrated that the true signals are still contained in the spoofed data. They have been located by using re-acquisition calculations that employ long coherent integration intervals in order to counteract the jamming effects of the high powered spoofer signals. The presence of the true navigation data bits on the spoofed signals aids in the implementation of long coherent accumulations.

The current prototype system could be augmented to re-do its signal acquisition calculations after the detection of a spoofing attack. After re-acquiring additional signals, it would need to do further hypothesis-testing-type calculations in order to make a definitive determination about which signals were the true ones.

G. Reuse of Stored RF Data from WRoD Spoofing Tests

Most of the suggested developments above can be tested using the stored USRP raw RF samples from the WRoD spoofing tests reported in this paper. These data contain a rich set of spoofing attack scenarios that will provide useful test cases for evaluating new algorithm designs and tunings. The needed testing can be done by re-playing the samples directly into a GNSS software receiver as though they were being sampled in real-time from an actual pair of antennas and USRPs.

VII. SUMMARY AND CONCLUSIONS

A new prototype GNSS spoofing detection system has been developed and tested using live-signal spoofing attacks. The system detects spoofing by using differences in signal direction-of-arrival characteristics between the spoofed and non-spoofed cases as sensed by a pair of GNSS antennas. A spoofing detection statistic has been developed that equals the difference between the optimized values of the negative-log-likelihood cost functions for two data fitting problems. One problem fits the single-differenced beat carrier phases of multiple received signals to a spoofed model in which the fractional parts of these differences are identical -- in the absence of receiver noise -- because the spoofed signals all arrive from the same direction. The other problem fits the single-differenced carrier phases to a non-spoofed model. This second optimal data fitting problem is closely related to CDGPS attitude determination. One of its by-products is an estimate of the reference-coordinates direction vector between the system's two antennas. The simple difference of the two optimized cost functions equals a large positive number if there is no spoofing, but it equals a negative number if the signals are being spoofed. Monte-Carlo analysis of the probability distributions of this difference under the spoofed and non-spoofed assumptions indicates that it provides a powerful spoofing detection test with a low probability of false alarm.

A real-time version of this system has been implemented using USRPs and real-time software radio receivers, and it has been tested against live-signal spoofing attacks aboard a yacht that was cruising around Italy. Successful detections have been achieved in many spoofing attack scenarios, and detections can occur in as little as 0.4 sec or less. One scenario spoofed the yacht's GPS receiver into believing that it had veered off of a northwesterly course towards Venice in the Adriatic to a southwesterly course towards the coast of Libya, and at the incredible speed of 900 kts. The spoofing detector, however, warned the crew on the bridge about the attack before the yacht's spoofed position was 50 m away from its true position.

The live-signal tests revealed some challenges for this spoofing detection strategy. They occur primarily during the initial attack phase, before the spoofer has dragged the victim receiver to a wrong position or timing fix. If the spoofer power is not very much larger than that of the true signals, then beating occurs between the spoofed and true signals during this initial period. This beating can cause difficulties for the receiver tracking loops, making single-differenced carrier phase unavailable. Even when single-differenced phase is available, both the spoofed and non-spoofed models of this quantity can be inadequate for purposes of designing a reliable spoofing detection test.

In summary, this paper's new two-antenna spoofing detection system has generated promising real-time results against live-signal spoofing attacks, but further developments are needed in order to produce a sufficiently reliable detection system for all anticipated attack scenarios. The best defense will likely employ a multi-layered approach that uses the techniques described in this paper along with advanced RAIM techniques which detect additional signal anomalies that are characteristic of spoofing.

ACKNOWLEDGEMENTS

The authors thank the owner of the White Rose of Drachs for the loan of his vessel in order to conduct the live-signal GNSS spoofing detection tests reported in this paper. The crew of the White Rose aided and supported this project in many ways.

REFERENCES

[1] Brown, R.G., "Receiver Autonomous Integrity Monitoring," in *Global Positioning System: Theory and Applications, Vol. II*, B.W. Parkinson and J.J. Spilker, Jr., eds., American Institute of Aeronautics and Astronautics, (Washington, 1996), pp. 143-165.

[2] Humphreys, T.E., Ledvina, B.M., Psiaki, M.L., O'Hanlon, B., and Kintner, P.M., Jr., "Assessing the Spoofing Threat: Development of a Portable GPS Civilian Spoofer," *Proc. ION GNSS 2008*, Sept. 16-19, 2008, Savannah, GA.

[3] Humphreys, T.E., Kintner, P.M., Jr., Psiaki, M.L., Ledvina, B.M., and O'Hanlon, B.W., "Assessing the Spoofing Threat," *GPS World*, Vol. 20, No. 1, Jan. 2009, pp. 28-38.

[4] Rawnsley, A., "Iran's Alleged Drone Hack: Tough, but Possible," *Wired*, available online at http://www.wired.com/dangerroom/2011/12/iran-drone-hack-gps/, Dec. 2011.

[5] Shepard, D.P., Bhatti, J.A., and Humphreys, T.E., "Drone Hack: Spoofing Attack Demonstration on a Civilian Unmanned Aerial Vehicle," *GPS World*, Vol. 23, No. 8, Aug. 2012, pp. 30-33.

[6] Kerns, A.J., Shepard, D.P., Bhatti, J.A., and Humphreys, T.E., "Unmanned Aircraft Capture and Control via GPS Spoofing," *Journal of Field Robotics*, Vol. 31, No. 4, 2014, pp. 617–636.

[7] Clark, L., "Students Push £50m Super Yacht Off Course using GPS Spoofing," *Wired*, available on-line at http://www.wired.co.uk/news/archive/2013-07/30/yacht-gps-hijack, July 2013.

[8] Zaragoza, S., and Zumalt, E., "Spoofing a Superyacht at Sea," Cockrell School of Engineering, Univ. of Texas/Austin, available on-line at

http://www.utexas.edu/know/2013/07/30/spoofing-a-superyacht-at-sea/, July 2013.

[9] Hartman, R.G., "Spoofing Detection System for a Satellite Positioning System," U.S. Patent No. 5,557,284, Sept. 1996.

[10] Montgomery, P.Y., Humphreys, T.E., and Ledvina, B.M., "A Multi-Antenna Defense: Receiver-Autonomous GPS Spoofing Detection," *Inside GNSS*, Vol. 4, No. 2, March/April 2009, pp. 40–46.

[11] Lo, S., De Lorenzo, D., Enge, P., Akos, D., and Bradley, P., "Signal Authentication, A Secure Civil GNSS for Today," *Inside GNSS*, Vol. 4, No. 5, Sept./Oct. 2009, pp. 30-39.

[12] Nielsen, J., Broumandan, A., and Lachapelle, G., "Method and System for Detecting GNSS Spoofing Signals," U.S. Patent No. 7,952,519 B1, May 2011.

[13] Levin, P., De Lorenzo, D.S., Enge, P.K., and Lo, S.C., "Authenticating a Signal Based on an Unknown Component Thereof," U.S. Patent No. 7,969,354 B2, June 2011.

[14] Dovis, F., Chen, X., Cavaleri, A., Ali, K., and Pini, M., "Detection of Spoofing Threats by Means of Signal Parameters Estimation," *Proc. ION GNSS 2011*, Sept. 20-23, 2011, Portland, OR, pp. 416-421.

[15] Bardout, Y., "Authentication of GNSS Position: An Assessment of Spoofing Detection Methods," *Proc. ION GNSS 2011*, Sept. 20-23, 2011, Portland, OR, pp. 436-446.

[16] Pini, M., Fantino, M., Cavaleri, A., Ugazio, S., and Lo Presti, L., "Signal Quality Monitoring Applied to Spoofing Detection," *Proc. ION GNSS 2011*, Sept. 20-23, 2011, Portland, OR, pp. 1888-1896.

[17] Shepard, D.P., and Humphreys, T.E., "Characterization of Receiver Response to a Spoofing Attacks," *Proc. ION GNSS 2011*, Sept. 20-23, 2011, Portland, OR, pp. 2608-2618.

[18] Psiaki, M.L., O'Hanlon, B.W., Bhatti, J.A., Shepard, D.P., and Humphreys, T.E., "Civilian GPS Spoofing Detection based on Dual-Receiver Correlation of Military Signals," *Proc. ION GNSS 2011*, Sept. 20-23, 2011, Portland, OR, pp. 2619-2645.

[19] Wesson, K.D., Shepard, D.P., Bhatti, J.A., and Humphreys, T.E., "An Evaluation of the Vestigial Signal Defense for Civil GPS Anti-Spoofing," *Proc. ION GNSS 2011*, Sept. 20-23, 2011, Portland, OR, pp. 2646-2656.

[20] Wesson, K.D., Rothlisberger, M.P., and Humphreys, T.E., "A Proposed Navigation Message Authentication Implementation for Civil GPS Anti-Spoofing," *Proc. ION GNSS 2011*, Sept. 20-23, 2011, Portland, OR, pp. 3129-3140.

[21] Nielsen, J., Broumandan, A., and Lachapelle, G., "GNSS Spoofing Detection for Single Antenna Handheld Receivers," *Navigation*, Vol. 58, No. 4, Winter 2011, pp. 335-344.

[22] Jafarnia-Jahromi, A., Lin, T., Broumandan, A., Nielsen, J., and Lachapelle, G., "Detection and Mitigation of Spoofing Attacks on a Vector Based Tracking GPS Receiver," *Proc. 2012 International Technical Meeting of the ION*, Jan. 30-Feb. 1, 2012, Newport Beach, CA, pp. 790-800.

[23] Broumandan, A., Jafarnia-Jahromi, A., Dehghanian, V., Nielsen, J., and Lachapelle, G., "GNSS Spoofing Detection in Handheld Receivers Based on Signal Spatial Correlation," *Proc. IEEE/ION PLANS 2012*, April 24-26, 2012, Myrtle Beach, SC, pp. 479-487.

[24] Daneshmand, S., Jafarnia-Jahromi, A., Broumandon, A., and Lachapelle, G., "A Low-Complexity GPS Anti-Spoofing Method Using a Multi-Antenna Array," *Proc. ION GNSS 2012*, Sept. 17-21, 2012, Nashville, TN, pp. 1233-1243.

[25] Trinkle, M., Zhang, Z., Li, H., and Dimitrovski, A., "GPS Anti-Spoofing Techniques for Smart Grid Applications," *Proc. ION GNSS 2012*, Sept. 17-21, 2012, Nashville, TN, pp. 1270-1278.

[26] Dehghanian, V., Nielsen, J., and Lachapelle, G., "GNSS Spoofing Detection based on Receiver C/N0 Estimates," *Proc. ION GNSS 2012*, Sept. 17-21, 2012, Nashville, TN, pp. 2875-2884.

[27] Meurer, M., Konovaltsev, A., Cuntz, M., and Hättich, C., "Robust Joint Multi-Antenna Spoofing Detection and Attitude Estimation using Direction Assisted Multiple Hypotheses RAIM," *Proc. ION GNSS 2012*, Sept. 17-21, 2012, Nashville, TN, pp. 3007-3016.

[28] Humphreys, T., Bhatti, J., Shepard, D., and Wesson, K., "The Texas Spoofing Test Battery: Toward a Standard for Evaluating GPS Signal Authentication Techniques," *Proc. ION GNSS 2012*, Sept. 17-21, 2012, Nashville, TN, pp. 3569-3583.

[29] Shepard, D.P., Bhatti, J.A., Humphreys, T.E., and Fansler, A.A., "Evaluation of Smart Grid and Civilian UAV Vulnerability to GPS Spoofing Attacks," *Proc. ION GNSS 2012*, Sept. 17-21, 2012, Nashville, TN, pp. 3591-3605.

[30] Wesson, K., Rothlisberger, M., and Humphreys, T., "Practical Cryptographic Civil GPS Signal Authentication", *Navigation*, Vol. 59, No. 3, Fall 2012, pp. 177-193.

[31] Daneshmand, S., Jafarnia-Jahromi, A., Broumandan, A., and Lachapelle, G. "Low Complexity Spoofing Mitigation," *GPS World*, Vol. 22, No. 12, Dec. 2012, pp. 44-46.

[32] Akos, D.M., "Who's Afraid of the Spoofer? GPS/GNSS Spoofing Detection via Automatic Gain Control (AGC)", *Navigation*, Vol. 59, No. 4, Winter 2012, pp. 281-290.

[33] Konovaltsev, A., Cuntz, M., Haettich, C., and Meurer, M., "Performance Analysis of Joint Multi-Antenna Spoofing Detection and Attitude Estimation," *Proc. 2013 International Technical*

Meeting of the ION, Jan. 29-27, 2013, San Diego, CA, pp. 864-872.

34 Swaszek, P.F., Hartnett, R.J., Kempe, M.V., and Johnson, G.W., "Analysis of a Simple, Multi-Receiver GPS Spoof Detector," *Proc. 2013 International Technical Meeting of the ION*, Jan. 29-27, 2013, San Diego, CA, pp. 884-892.

35 Humphreys, T.E, "Detection Strategy for Cryptographic GNSS Anti-Spoofing," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 49, No. 2, April 2013, pp. 1073-1090.

36 Psiaki, M.L., Powell, S.P., and O'Hanlon, B.W., "GNSS Spoofing Detection, Correlating Carrier Phase with Rapid Antenna Motion," *GPS World*, Vol. 24, No. 6, June 2013, pp. 53-58.

37 Moon, G.B., Im, S.H., and Jee, G.-I., "A Civil GPS Anti-Spoofing and Recovering Method Using Multiple Tracking Loops and an Adaptive Filter Technique," *Proc. ION GNSS+ 2013*, Sept. 17-20, 2013, Nashville, TN, pp. 2916-2920.

38 Swaszek, P.F., and Hartnett, R.J., "Spoof Detection Using Multiple COTS Receivers in Safety Critical Applications," *Proc. ION GNSS+ 2013*, Sept. 17-20, 2013, Nashville, TN, pp. 2921-2930.

39 Dehghanian, V., and Nielsen, J., "GNSS Spoofing Detection based on a Sequence of RSS Measurements," *Proc. ION GNSS+ 2013*, Sept. 17-20, 2013, Nashville, TN, pp. 2931-2936.

40 Konovaltsev, A., Cuntz, M., Haettich, C., and Meurer, M., "Autonomous Spoofing Detection and Mitigation in a GNSS Receiver with an Adaptive Antenna Array," *Proc. ION GNSS+ 2013*, Sept. 17-20, 2013, Nashville, TN, pp. 2937-2948.

41 Psiaki, M.L., Powell, S.P., and O'Hanlon, B.W., "GNSS Spoofing Detection Using High-Frequency Antenna Motion and Carrier-Phase Data," *Proc. ION GNSS+ 2013*, Sept. 17-20, 2013, Nashville, TN, pp. 2949-2991.

42 Turner, M., Chambers, A., Mak, E., Aguado, L.E., Wales, B., and Dumville, M., "PROSPA: Open Service Authentication," *Proc. ION GNSS+ 2013*, Sept. 17-20, 2013, Nashville, TN, pp. 2992-2996.

43 Nielsen, J., Dehghanian, V., and Dawar, N., "GNSS Spoofing Detection based on Particle Filtering," *Proc. ION GNSS+ 2013*, Sept. 17-20, 2013, Nashville, TN, pp. 2997-3005.

44 Psiaki, M.L., O'Hanlon, B.W., Bhatti, J.A., Shepard, D.P., and Humphreys, T.E., "GPS Spoofing Detection via Dual-Receiver Correlation of Military Signals," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 49, No. 4, Oct. 2013, pp. 2250-2267.

45 Wesson, K.D., Evans, B.L., and Humphreys, T.E., "A Combined Symmetric Difference and Power Monitoring GNSS Anti-Spoofing Technique," *Proc. IEEE Global Conference on Signal and Information Processing* (*GlobalSIP*), Dec. 3-5, 2013, Austin, TX.

46 O'Hanlon, B.W., Psiaki, M.L., Bhatti, J.A., Shepard, D.P., and Humphreys, T.E., "Real-Time GPS Spoofing Detection via Correlation of Encrypted Signals" *Navigation*, Vol. 60, No. 4, Winter 2013, pp. 267-278.

47 Li, Z., and Gebre-Egziabher, D., "Performance Analysis of a Civilian GPS Position Authentication System", *Navigation*, Vol. 60, No. 4, Winter 2013, pp. 246-265.

48 Swaszek, P.F., and Hartnett, R.J., "A Multiple COTS Receiver GNSS Spoof Detector -- Extensions," *Proc. 2014 International Technical Meeting of the ION*, Jan. 27-29, 2014, San Diego, CA, pp. 316-326.

49 Jafarnia-Jahromi, A., Broumandan, A., Nielsen, J., and Lachapelle, G., "Pre-Despreading Authenticity Verification for GPS L1 C/A Signals", *Navigation*, Vol. 61, No. 1, Spring 2014, pp. 1-11.

50 Psiaki, M.L., "Spoofing Detection for Civilian GNSS Signals," U.S. Patent No. 8,712,051 B2, April 2014.

51 Hwang, P.Y., and McGraw, G.A., "Receiver Autonomous Signal Authentication (RASA) Based on Clock Stability Analysis," *Proc. IEEE/ION PLANS 2014*, May 5-8, 2014, Monterey, CA, May 2014, pp. 270-281.

52 Yuan, D., Li, H., and Lu, M., "A Method for GNSS Spoofing Detection Based on Sequential Probability Ratio Test," *Proc. IEEE/ION PLANS 2014*, May 5-8, 2014, Monterey, CA, pp. 351-358.

53 Khanafseh, S., Roshan, N., Langel, S., Chan, F.-C., Joerger, M., and Pervan, B., "GPS Spoofing Detection Using RAIM with INS Coupling," *Proc. IEEE/ION PLANS 2014*, May 5-8, 2014, Monterey, CA, pp. 1232-1239.

54 Jovanovic, A., Botteron, C., and Fariné, P.-A., "Multi-test Detection and Protection Algorithm Against Spoofing Attacks on GNSS Receivers," *Proc. IEEE/ION PLANS 2014*, May 5-8, 2014, Monterey, CA, pp. 1258-1271.

55 Kerns, A.J., Wesson, K.D., and Humphreys, T.E., "A Blueprint for Civil GPS Navigation Message Authentication," *Proc. IEEE/ION PLANS 2014*, May 5-8, 2014, Monterey, CA, pp. 262-269.

56 Giorgi, G., Teunissen, P.J.G., Verhagen, S., and Buist, P.J., "Instantaneous Ambiguity Resolution in Global-Navigation-Satellite-System-Based Attitude Determination Applications: A Multivariate Constrained Approach", *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 1, Jan.-Feb. 2012, pp. 51-67.

57 Psiaki, M.L., "Batch Algorithm for Global-Positioning-System Attitude Determination and Integer Ambiguity Resolution", *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 5, Sept.-Oct. 2006, pp. 1070-1079.

[58] Gill, P.E., Murray, W., and Wright, M.H., *Practical Optimization*, Academic Press, (New York, 1981), pp. 37-40.

[59] Poor, H.V., *An Introduction to Signal Detection and Estimation*, Springer-Verlag, (New York, 1988), pp. 7-195.

[60] Van Trees, H.L., Bell, K.L., and Tian, Z., *Detection, Estimation, and Modulation Theory, 2nd Ed.*, J. Wiley & Sons, (Hoboken, NJ, 2013), p. 359.

[61] Psiaki, M.L., Ertan, T., O'Hanlon, B.W., and Powell, S.P., "GNSS Multipath Mitigation using High-Frequency Antenna Motion," submitted to *Navigation*, in review. An earlier version appears in *Proc. ION GNSS+ 2013*, Sept. 17-20, 2013, Nashville, TN, pp. 154-175.