Title
GNSS Spoofing Detection: An Approach for Ground Vehicles Using Carrier-Phase and Inertial Measurement Data

Journal
GPS World Magazine, Volume 34, Issue 2

Authors
Clements, Z
Yoder, James E.
Humphreys, Todd E.

Publication Date
2023-02

GPS WORLD
GNSS POSITIONING NAVIGATION TIMING
THE INDUSTRY'S MOST TRUSTED TECHNICAL RESOURCE SINCE 1990

INNOVATION
GNSS Spoofing Detection

ALIGNING BRICKS AND MODELS

Surveying for Architecture, Engineering & Construction

+ WHO RUNS GPS?

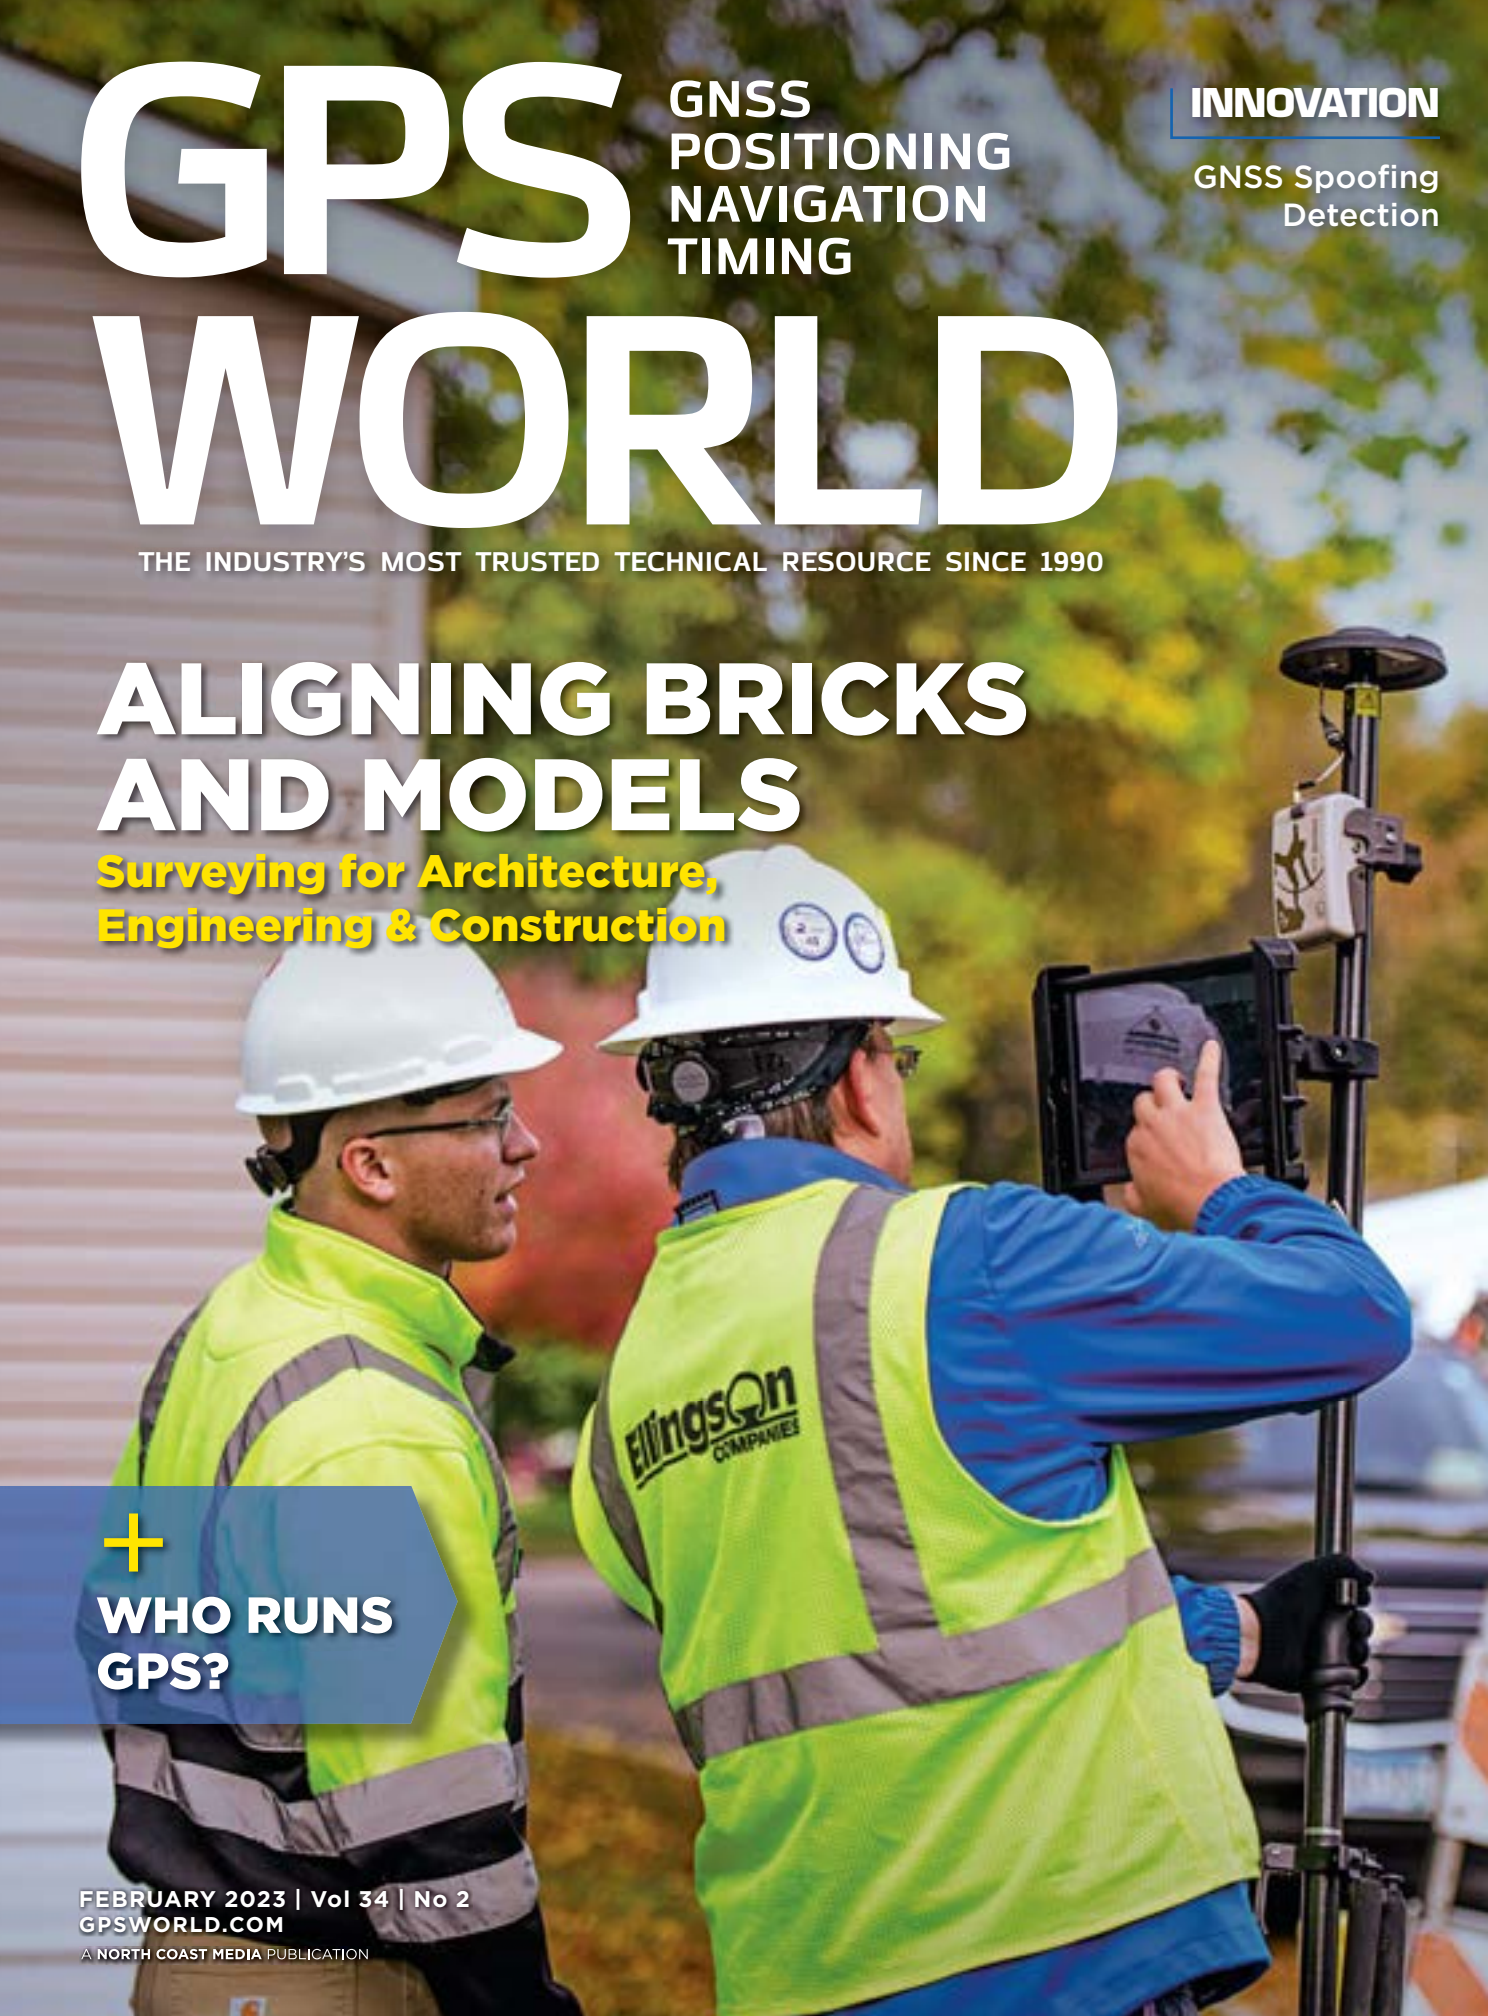

FEBRUARY 2023 | Vol 34 | No 2
GPSWORLD.COM
A NORTH COAST MEDIA PUBLICATION

FIGURE 1 The UT RNL has developed a multi-modal ground-vehicle-mounted integrated perception platform call the Sensorium. It houses three different types of IMU, two triple-frequency GNSS antennas, three radar sensors and two cameras.

All figures provided by the authors.

GNSS SPOOFING DETECTION

An Approach for Ground Vehicles Using Carrier-Phase and Inertial Measurement Data

BY ZACHARY CLEMENTS, JAMES E. YODER AND TODD E. HUMPHREYS

The combination of easily accessible low-cost GNSS spoofers and the emergence of increasingly automated GNSS-reliant ground vehicles prompts a need for fast and reliable GNSS spoofing detection. To underscore this point, Regulus Cyber, an Israeli cybersecurity company, recently spoofed a Tesla Model 3 on autopilot mode, causing the vehicle to suddenly slow and unexpectedly veer off the main road.

Among GNSS signal authentication techniques, signal-quality monitoring (SQM) and multi-antenna could be considered for implementation on ground vehicles. However, SQM tends to perform poorly on dynamic platforms in urban areas where strong multipath and in-band noise are common, and multi-antenna spoofing detection techniques, while effective, are disfavored by automotive manufacturers seeking to reduce vehicle cost and aerodynamic drag. Thus, there is a need for a single-antenna GNSS spoofing detection technique that performs well on ground vehicles, despite the adverse signal-propagation conditions in an urban environment.

In a concurrent trend, increasingly automated ground vehicles demand ever-stricter lateral positioning to ensure safety of operation. An influential study calls for lateral positioning better than 20 centimeters on freeways and better than 10 centimeters on local streets (both at a 95% probability level). Such stringent requirements can be met by referencing lidar and camera measurements to a local high-definition map, but poor weather (heavy rain, dense fog or snowy whiteout) can render this technique unavailable.

On the other hand, progress in precise (decimeter-level) GNSS-based ground vehicle positioning, which is impervious to poor weather, has demonstrated surprisingly high (above 97%) solution availability in urban areas. This technique is based on carrier-phase differential GNSS (CDGNSS) positioning, which exploits GNSS carrier-phase measurements having millimeter-level precision but integer-wavelength ambiguities.

Key to our promising results is the tight coupling of CDGNSS and inertial measurement unit (IMU) data, without which high-accuracy CDGNSS solution availability is significantly reduced due to pervasive signal blockage and multipath in urban areas. Tight coupling brings millimeter-precise GNSS carrier-phase measurements into correspondence with high-sensitivity and high-frequency inertial sensing. Our particular estimation architecture incorporates inertial sensing via model replacement, in which the estimator's propagation step relies on bias-compensated acceleration and angular rate measurements from the IMU instead of a vehicle dynamics model. As a consequence, at each measurement update, an a priori antenna position is available whose delta from the previous measurement update accounts for all vehicle motion sensed by the IMU, including small-amplitude high-frequency motion caused by road irregularities. Remarkably, when tracking authentic GNSS signals in a clean (open-sky) environment, the GNSS carrier-phase predicted by the a priori antenna position and the actual measured carrier phase agree to within millimeters.

The research described in this article pursues a novel GNSS spoofing-detection technique based on a simple but consequential observation: it is practically impossible for a spoofer to create a false ensemble of GNSS signals whose

carrier-phase variations, when received through the antenna of a target ground vehicle, track the phase values predicted by inertial sensing. In other words, antenna motion caused by factors such as road irregularities or rapid braking or steering is sensed with high fidelity by an onboard IMU but is unpredictable at the sub-centimeter-level by a would-be spoofer. Therefore, the differences between IMU-predicted and measured carrier-phase values offer the basis for an exquisitely sensitive GNSS spoofing-detection statistic. What is more, such carrier-phase fixed-ambiguity residual cost is generated as a byproduct of tightly coupled inertial-CDGNSS vehicle position estimation.

Two difficulties complicate the use of fixed-ambiguity residual cost for spoofing detection. First is the integer-ambiguous nature of the carrier-phase measurement, which causes the post-integer-fix residual cost to equal not the difference between the measured and predicted carrier phases (as would be the case for a typical residual), but rather modulo an integer number of carrier wavelengths. Such integer folding complicates development of a probability distribution for a detection test statistic based on carrier-phase fixed-ambiguity residual cost.

Second, the severe signal multipath conditions in urban areas create thick tails in any detection statistic based on carrier-phase measurements. Setting a detection threshold high enough to avoid false spoofing alarms caused by mere multipath could render the detection test insensitive to dangerous forms of spoofing. Reducing false alarms by accurately modeling the effect of a particular urban multipath environment on the detection statistic would be a Sisyphean undertaking, requiring exceptionally accurate up-to-date 3D models of the urban landscape, including materials properties.

Our work takes an empirical approach to these difficulties. It does not attempt to develop a theoretical model to delineate the effects of integer folding or multipath on its proposed carrier-phase fixed-ambiguity residual cost-based detection statistic. Rather, it develops null-hypothesis empirical distributions for the statistic in both shallow and deep urban areas, and uses these distributions to demonstrate that high-sensitivity spoofing detection is possible despite integer folding and urban multipath.

MEASUREMENT MODEL

The full formulation of the measurement model for the tightly coupled GNSS-IMU estimator on which our spoofing detection technique is based is presented in a paper accepted for publication in the Institute of Navigation's journal *NAVIGATION*. Below are key developments.

The estimator ingests pairs of double-difference (DD) GNSS observables at each GNSS measurement epoch, with each pair composed of a pseudorange and a carrier-phase measurement. After linearizing about the a priori state estimate, a measurement model for the innovations vector can be expressed by setting the difference between the measurement vector and its modeled value based on the a priori state estimate equal to the corresponding Jacobians (matrices of partial derivatives), the state estimate error vector, and the integer ambiguity vector. A short-baseline

INNOVATION INSIGHTS

BY RICHARD B. LANGLEY

WHAT IS CARRIER PHASE? The obvious answer is: the phase of the carrier. But this is not helpful if you don't know what a carrier is. A carrier is basically a harmonic electromagnetic wave — a pure continuous sinusoidal wave with a single constant frequency and amplitude.

Such a wave has limited uses. However, if we modulate or change the characteristics of the wave in some way, then the wave can carry information. Changing the amplitude by using a voice or music audio signal is amplitude modulation as used for AM radio. Instead, one could modulate a carrier by changing its instantaneous frequency, which is frequency modulation or FM and is used for high-fidelity broadcasting. Yet another way to modulate a carrier is to change the instantaneous phase of the carrier, and that is how GNSS works.

GNSS carriers are phase-modulated by pseudorandom noise (PRN) codes and navigation messages. A GNSS receiver uses the PRN codes to produce the pseudorange observable with a precision in the tens of decimeter range. This is the most common observable for GNSS positioning.

But by stripping away the modulation of the received GNSS signals, the receiver can measure the phase of the underlying carrier. Changes in carrier phase over time reflect the change in the (pseudo) range but are about two orders of magnitude more precise. One problem with carrier-phase measurements is that they have an initial cycle ambiguity that must be resolved, preferentially fixed to the correct integer value, before they can be used for positioning, but this can be achieved without too much difficulty.

While fixing the ambiguity of carrier-phase measurements might be considered a nuisance in GNSS positioning, it can help detect spoofing of GNSS signals where some other techniques might fall short. In this "Innovation" column, we look at how carrier-phase measurements combined with those from an inertial measurement unit can guard against a deliberate attack on an automated ground vehicle — something that cannot be discounted in our world these days.

regime is assumed for the DD measurements, which implies that ionospheric, tropospheric, ephemeris and clock errors are cancelled in the double differencing.

The CDGNSS measurement update of the tightly coupled GNSS-IMU estimator can be cast in square-root form for greater numerical robustness and algorithmic clarity. The measurement update can be defined as the process of finding the state estimate error vector and the integer ambiguity vector to minimize the cost function, a term borrowed from economics theory, which is essentially the sum of the squared errors. The cost function can be decomposed into three terms: a term that can be zeroed for any value of the integer ambiguity vector; a term corresponding to the residual cost of enforcing the integer constraint; and the irreducible cost that can be shown to be equivalent to the normalized innovations associated with the DD pseudorange measurements. We normalize the vector cost components by using square-root information matrices based on Cholesky factorization followed by the cost decomposition using QR factorization. The equations are then solved to provide float (real number) values for the state estimate error vector and the integer ambiguity vector. Subsequently, the fixed solution is found via an integer least squares (ILS) solver.

TEST STATISTIC

Key to our spoofing detection statistic is the integer-fixed carrier-phase residual cost value, which also can be thought of as the ILS solution cost. This is equivalent to the second term of the decomposed cost function stated earlier. This is small whenever the carrier-phase measurements are consistent with the prior state estimate and the pseudorange measurements, and with the assumption of integer-valued carrier-phase ambiguities. It is one of several acceptance test statistics used to decide whether the fixed solution is correct with high probability. It has been incorporated, for example, in a statistic used to detect carrier cycle slips. It can be similarly used to detect false integer fixes, just as with other integer aperture acceptance test statistics, or the lingering effects of conditioning the real-valued part of the state on a previous false fix.

Furthermore, the test statistic provides a highly sensitive statistic for spoofing detection. When no spoofing is present, there is tight agreement between the IMU-propagated a priori state estimate and GNSS data resulting in a small statistic value. If the vehicle hits a bump in the road, the GNSS antenna phase center will rise by a few centimeters, and the inertial sensor will detect a corresponding acceleration, which will get propagated through to the a priori state. On the other hand, when spoofing is present, a discrepancy between inertial and GNSS data will arise at the carrier-phase level, leading to the statistic being larger than usual.

A windowed sum of test statistic values offers even greater sensitivity to false-fix events at the expense of a longer time-to-detect. To detect spoofing in the tests reported in this article, we used the windowed fixed-ambiguity residual cost (WFARC). This is calculated over a moving window of fixed length of past GNSS measurement epochs. We used a window length of 10 past GNSS measurement epochs (amounting to a window of 2 seconds).

If the filter is consistent and the integer ambiguities are correctly resolved, then WFARC should be approximately χ^2-distributed with the degrees of freedom related to the number of DD carrier-phase measurements. This distribution is approximate due to the "integer-folding" effect: large phase residuals are not possible because of integer-cycle phase wrapping. A statistical consistency test can be performed by choosing a desired false-alarm rate and declaring a false fix if WFARC is greater than a specified threshold calculated by evaluating the inverse cumulative distribution function of χ^2 at the false-alarm rate.

DATA COLLECTION

Data was gathered on the University of Texas (UT) Radionavigation Laboratory (RNL) Sensorium, an integrated platform for automated and connected vehicle perception research. It is equipped with multiple radars, IMUs, GNSS receivers and a lidar, as shown in **FIGURE 1**. With the Sensorium, the RNL produced a public benchmark dataset collected in the dense urban center of the city of Austin called TEX-CUP for evaluating multi-sensor GNSS-based urban positioning algorithms. The data captured includes a diverse set of multipath environments (open-sky, shallow urban and deep urban). The TEX-CUP dataset provides raw wideband intermediate frequency (IF) GNSS data with tightly synchronized raw measurements from multiple IMUs and a stereoscopic camera unit, as well as truth positioning data. This allows researchers to develop algorithms using any subset of the sensor measurements and compare their results with the true position.

For our analysis, only the raw GNSS IF samples from the primary antenna and inertial data were considered. Two-bit-quantized IF samples were captured at the Sensorium and at the reference station through the RadioLynx, a low-cost L1+L2 GNSS front end with a 5 MHz bandwidth at each frequency, and were processed with the RNL's GRID software-defined radio (SDR). The system's performance was separately evaluated using inertial data from each of the Sensorium's two MEMS inertial sensors. The first, a LORD MicroStrain 3DM-GX5-25, is an industrial-grade sensor. The second, a Bosch BMX055, is a surface-mount consumer-grade sensor.

TEX-CUP provides ground truth data for the vehicle position and orientation. The post-processed solution is accurate to better than 10 centimeters throughout the dataset. The effectiveness of the developed spoofing detection method is evaluated with the dataset subsets.

SPOOFING METHODOLOGY

The total signal at the victim receiver antenna is the sum

of the authentic signal, the spoofed signal and the received noise. Under a challenging spoofing attack, the spoofed signal contains a perfect null of the authentic signal and the received noise, which is entirely naturally generated — that is, not introduced by the spoofer.

Physical-Layer Spoofing. To artificially simulate a spoofing attack over-the-air, cable injection and digital signal injection spoofing were considered. Over-the-air attacks are possible, but are not authorized in urban areas. A cable injection attack would be permissible for a live experiment in an urban area, and digital signal combining, is a powerful after-the-fact spoofing technique. But in both cases it is challenging to explore a worst-case spoofing attack in which the authentic signals are entirely nulled by an antipodal spoofing signal. Experience with ds7 and ds8 from the Texas Spoofing Test Battery (**http://radionavlab.ae.utexas.edu/texbat**) revealed that such antipodal spoofing is difficult to maintain under even static laboratory conditions. The remnant authentic signal from an unsophisticated and imperfect spoofing attack sullies the test statistic, making detection too easy and leading to an overly optimistic performance assessment.

In short, physical-layer spoofing is challenging to conduct in such a way as to present a convincing worst-case spoofing attack to our detector.

Observation-Domain Spoofing. It is important to evaluate spoofing detection techniques on a worst-case spoofing attack, with the idea being that if the proposed detection strategy is effective on the worst-case scenario, it is even more effective on weaker attacks. Accordingly, we adopt *observation-domain spoofing*. The spoofing in the observation domain is advantageous because the authentic signal is inherently nulled, presenting a subtle attack.

The first method of implementing observation-domain spoofing is *position offset spoofing*. With position offset spoofing, a position offset is added to the authentic measured position to generate a spoofed position. This is accomplished by altering the pseudorange and carrier-phase measurements from each satellite so that they correspond to the spoofed position with the desired additive position offset.

The second method of implementing observation-domain spoofing is *timestamp spoofing*. With timestamp spoofing, the measurements at a particular time are reassigned to have an alternate measurement timestamp. The authentic observables from time $t+\delta t(t)$ are fed to the estimator as if they had occurred at time t. The timestamp-shifted observables are adjusted to account for the transmitting spacecraft's orbital motion and clock evolution over the interval from t to $t+\delta t(t)$.

In position offset spoofing, all vehicle motion reflected in the authentic carrier-phase observation is also present in the spoofed observation. This includes all high-frequency motion due to the road irregularities and other minor movements. A detection technique designed to detect small-amplitude, high-frequency discrepancies in carrier-phase measurements via the WFARC would not actually see such discrepancies unless the change in carrier phase due to the position offset also included simulated high-frequency content.

By contrast, timestamp spoofing borrows spoofed carrier-phase and pseudorange measurements from a different time instant, ensuring that high-frequency variations in these quantities will be different from those predicted by the a priori state based on IMU propagation. This is more representative of an actual spoofing attack scenario in which the attacker cannot predict the high-frequency vehicle motion. Moreover, by reducing the timestamp shift $\delta t(t)$, one can realize ever-subtler attacks that are increasingly hard to detect, allowing exploration of worst-case-for-detectability spoofing.

Thus, timestamp spoofing is representative of a case in which a well-financed attacker is able to place a single-satellite-full-single-ensemble spoofer capable of full authentic-signal nulling along the line-of-sight from the target vehicle to each overhead GNSS satellite.

RESULTS

We analyzed the proposed test statistics in both the non-spoofing case and against a worst-case attack, and present results with both the industrial- and consumer-grade IMUs.

Characterization of the Null Hypothesis. Our spoofing detector is premised on a hypothesis test between statistical models for the authentic and counterfeit GNSS signals. The statistics of the null hypothesis (no spoofing detected) must be fully characterized so that a statistical baseline is established, against which carrier-phase errors induced by spoofing in the same setting can be compared. The null hypothesis of dynamic ground vehicle scenarios includes natural effects such as blockage and multipath, which is the predominant source of error.

To analyze the null hypothesis, the WFARC was calculated in the nominal case through the entirety of the TEX-CUP dataset containing no spoofing. Because multipath is dependent on the surrounding environment, two categories were separately considered: shallow urban and deep urban. We separated measurements into these categories manually by identifying segments of the dataset where the vehicle resided in shallow urban and deep urban areas.

FIGURE 2 shows the complementary cumulative distribution function (CCDF) of the WFARC in shallow and deep urban environments for the nominal case with industrial- and consumer-grade IMUs. The test statistic in the deep urban case has a much longer tail, which is expected because of the extreme multipath and blockage in deep urban areas. The cyan line represents the largest value of the WFARC in the shallow urban environment and the purple line represents the largest value of the WFARC in the deep urban environment. These will be the thresholds used to detect spoofing. Because the test statistic in the null hypothesis is never larger than these values, it corresponds to having a false alarm probability of zero. A chi-

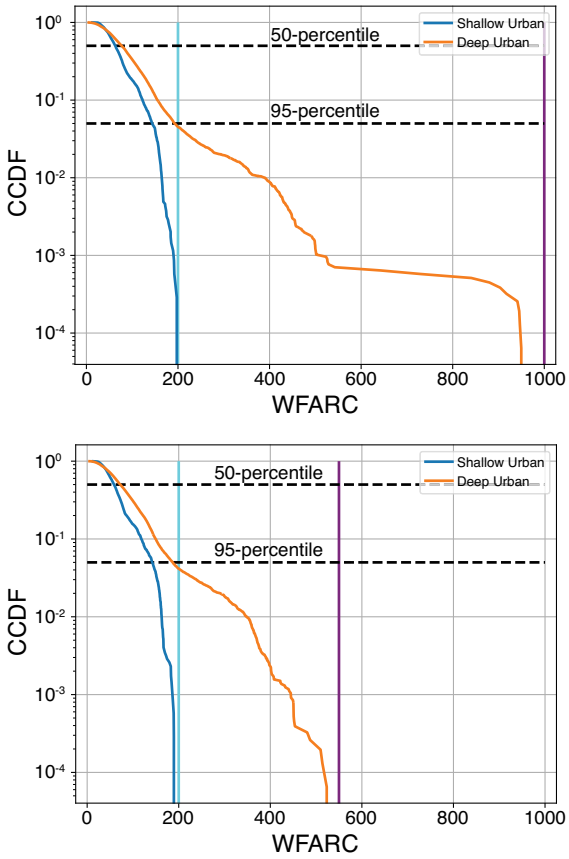

FIGURE 2 The complementary cumulative distribution function (CCDF) of the WFARC over the entire TEX-CUP dataset with the LORD MicroStrain 3DM-GX5-25 (industrial grade) IMU (top) and with the Bosch BMX055 (consumer grade) IMU (bottom).

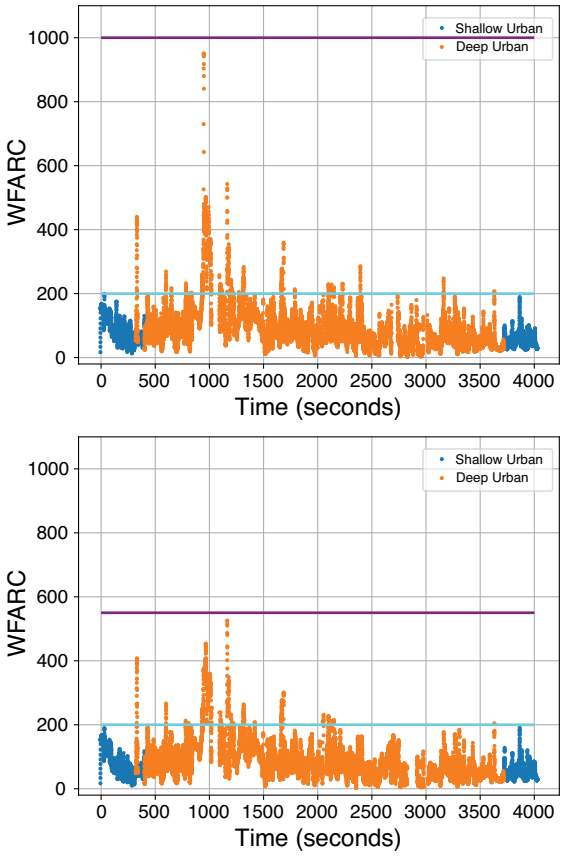

FIGURE 3 A time history of the WFARC over the entire TEX-CUP dataset with the industrial grade IMU (top) and the consumer grade IMU (bottom).

squared test can be used to lower these thresholds but comes at the cost of having a fixed false positive rate. **FIGURE 3** shows the time history of the WFARC over the TEX-CUP dataset.

It is important to note that the WFARC while using the consumer-grade IMU is generally smaller than the WFARC while using the industrial-grade IMU. This is expected because the consumer-grade IMU is of lesser quality, thus having more variance with each measurement. The a priori state estimate from IMU tight coupling has a larger uncertainty because the estimator has less confidence in the IMU measurements, leading corrupted measurements to be more believable. Once again, in the null hypothesis, spikes in the WFARC are caused by multipath and blockage.

Performance Against a Worst-Case Spoofing Attack. The following is an example of a worst-case spoofing attack in a shallow urban environment. In this scenario, the spoofing attack begins while the vehicle is stopped at a stoplight and continues as the vehicle begins to move. The WFARC in this scenario is shown in **FIGURE 4** with both industrial- and consumer-grade IMUs. The vehicle starts moving at the 163-second mark. The spoofing attack begins at the 163-second mark just before first movement and ends at the 175-second mark. As the vehicle begins to move, the position errors will grow gradually because the vehicle slowly begins to accelerate forward, inducing a position error. Three different time-shift attacks in the same scenario are shown in this figure. The shift of 0.15 seconds is the least subtle attack while the 0.05 second attack is the most challenging attack because the faults are much smaller. As the vehicle begins to move, the estimator recognizes inconsistencies between the spoofed GNSS measurements and the IMU because of the tight coupling. The rise in the WFARC above the thresholds shows this disagreement that is attributed to spoofing.

With the industrial-grade IMU and using the shallow urban threshold, all three time shifts spoofing attacks were identified within a second. The estimator knows that the IMU data are different than the GNSS measurements from the WFARC, much more than anything multipath would induce in the shallow urban environment. If the vehicle was in the deep urban environment, the 0.05-second shift spoofing attack would just be attributed to multipath. The sensitivity of the test is dependent on the multipath environment.

All three attacks were identified within two seconds while using the consumer-grade IMU. If the deep urban threshold

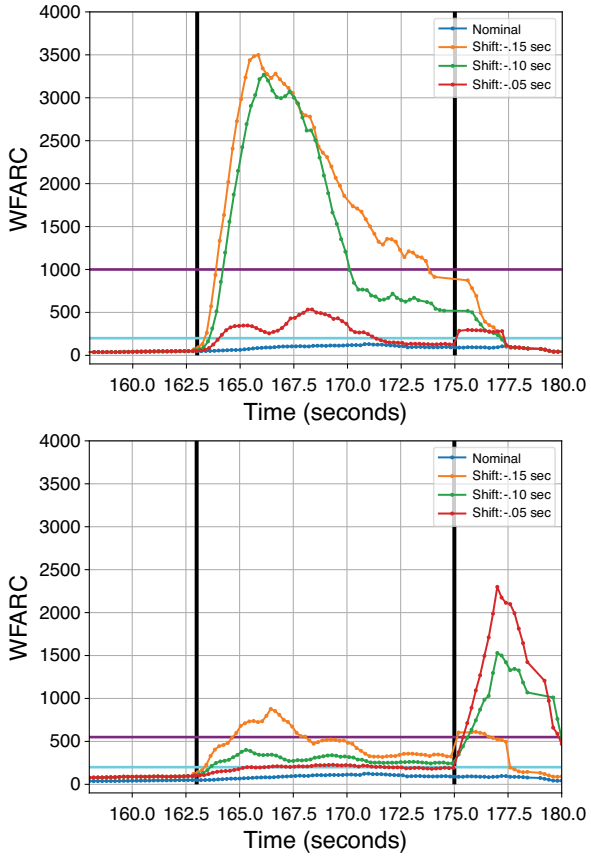

FIGURE 4 **WFARC during a worst-case spoofing attack in the shallow urban environment. The top plot is with the industrial grade IMU and the bottom plot is with the consumer grade IMU.**

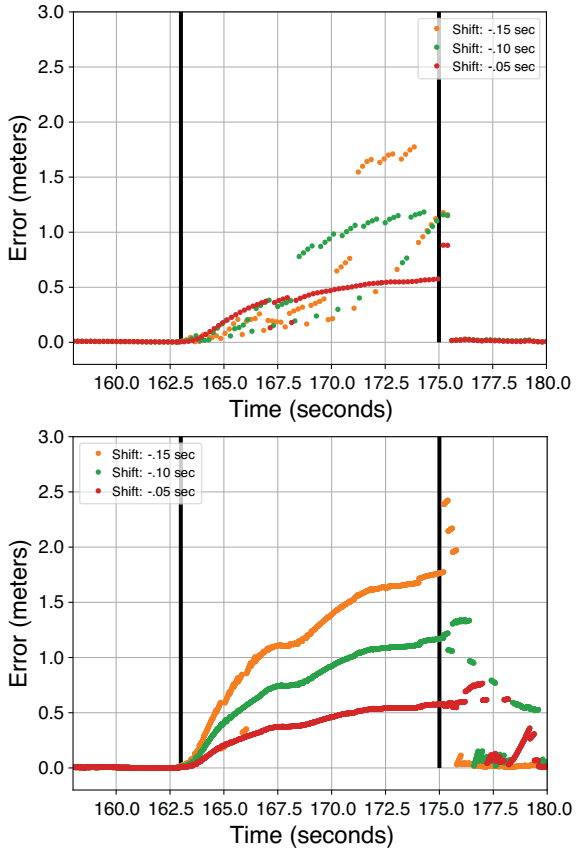

FIGURE 5 **The position errors induced from the different spoofing attacks. The top plot is with the industrial grade IMU and the bottom plot is with the consumer grade IMU.**

had been applied, only the least challenging attack would have been identified. In all cases, the WFARC is significantly smaller compared to when the industrial IMU is used. Once again, this is because the estimator has more confidence in the spoofed GNSS measurements than the lower quality IMU. Interestingly, there is a spike in the WFARC after the spoofing attack is over. This happens because the estimator is showing trauma from the spoofing attack — the abrupt return of the true GNSS measurements, which were significantly different from what the previously ingested spoofed measurements were predicting.

The corresponding position errors in each attack are shown in **FIGURE 5**. The worst-case attack (time shift of –0.05 seconds) only introduces a 0.5 meter offset over 10 seconds, indicative of an extremely subtle attack. Even the least subtle attack (time shift of –0.15 seconds) only introduces a 2-meter offset after 10 seconds, which is much more challenging than the attacks simulated in the related work.

FUTURE WORK

The results from our work are promising. It would be beneficial to collect even more data with the Sensorium to strengthen the empirical model.

ACKNOWLEDGMENTS

The work described in this article was supported in part by the U.S. Department of Transportation and by the Army Research Office. The views and conclusions are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. government. This article is based on the paper "Carrier-phase and IMU based GNSS Spoofing Detection for Ground Vehicles" presented at ION ITM 2022, the 2022 International Technical Meeting of the Institute of Navigation, Long Beach, California, Jan. 25–27, 2022.

ZACHARY CLEMENTS (BS, Electrical Engineering, Clemson University) is a graduate student in the Department of Aerospace Engineering, and Engineering Mechanics at The University of Texas (UT) at Austin, and a member of the UT Radionavigation Laboratory.

JAMES YODER (BS, Electrical Engineering, MS, Aerospace Engineering, UT Austin) is a guidance, navigation and control engineer at SpaceX and an alumnus of the UT Radionavigation Laboratory.

TODD HUMPHREYS (BS, MS, Electrical Engineering, Utah State University; PhD, Aerospace Engineering, Cornell University) is a professor in the Department of Aerospace Engineering and Engineering Mechanics at UT Austin, where he directs the Radionavigation Laboratory.