Providing Continuity and Integrity in the Presence of GNSS Spoofing

Fabian Rothmaier, Yu-Hsuan Chen, Sherman Lo, Juan Blanch, Todd Walter, *Stanford University*

BIOGRAPHY

Fabian Rothmaier is a PhD candidate at the GPS Laboratory at Stanford University. He received his B. Engr. degree from the University of Applied Sciences Bremen, Germany in 2015 and his M. Sc. degree from Stanford University in 2017.

Yu-Hsuan Chen is a research associate at the Stanford GPS Laboratory. He received his Ph.D. in electrical engineering from National Cheng Kung University, Taiwan.

Sherman Lo is a senior research engineer at the Stanford GPS Laboratory. He received his Ph.D. in Aeronautics and Astronautics from Stanford University in 2002. He has and continues to work on navigation robustness and safety, often supporting the FAA. He has conducted research on Loran, alternative navigation, SBAS, ARAIM, GNSS for railways and automobile. He also works on spoof and interference mitigation for navigation. He has published over 100 research papers and articles.

Juan Blanch is a Senior Research Engineer with Stanford University, where he works on the integrity algorithms for space-based augmentation systems and on receiver autonomous integrity monitoring. He received a Ph.D. in aeronautics and astronautics from Stanford University in 2003. He has received The Institute of Navigation (ION) Parkinson and Early Achievement awards.

Todd Walter received his Ph.D. in Applied Physics from Stanford University in 1993. He is a Research Professor in the Department of Aeronautics and Astronautics at Stanford University. His research focuses on implementing high-integrity air navigation systems. He has received the ION Thurlow and Kepler awards. He is also a fellow of the ION and has served as its president.

ABSTRACT

In this paper we attempt to provide continuity and integrity guarantees despite the presence of a spoofing attack. Using a software defined radio (SDR), we assign two channels per PRN allowing us to track both the authentic and spoofed signal. We employ a Multi-Hypothesis Extended Kalman Filter (MHEKF) to coast on inertial sensors during signal outages and identify the authentic signals once they become available. We finally leverage and modify RAIM integrity equations to cast protection levels that encompass the authentic and spoofed position solution for a guarantee on the user's location.
We present results both for highway driving data with artificially injected spoofing signals as well as results on the TEXBAT dataset. We demonstrate continuous navigation despite the presence of spoofing attacks and provide integrity through protection levels that account for the lift-off spoofing attacks in their threat model.

I. INTRODUCTION

GNSS has become the foundation of the position, velocity and time (pvt) solution of many safety of life applications. This success story has been fueled by the accuracy, availability, continuity and high levels of integrity offered by satellite navigation systems. Nowadays this performance is put in jeopardy by intentional and unintentional interference.
In this paper we consider targeted interference through counterfeit GNSS signals. Generally broadcasted by a malicious actor attempting to fool the victim's GNSS receiver, this attack on satellite navigation systems is known as spoofing. A rich body of research exists on GNSS spoofing detection, and promising results have been shown for defenses against different threat scenarios. Overviews of attack modes and common defense strategies are given by [1–3]. While there is no silver bullet among defense strategies, several approaches combining metrics have shows robust behavior detecting a wide range of threats [4, 5].

In this paper we consider the less thoroughly researched topic of GNSS spoofing mitigation. Mitigation is the attempt to resume the use of satellite navigation despite the presence of an attack. The goal is to exclude the compromised satellite signals from the navigation solution and continue the use of the authentic signals.

The idea to exclude unwanted signals from the navigation solution goes at least back to the failure detection and isolation suggested in [6], nowadays generally implemented within the scope of Receiver Autonomous Integrity Monitoring (RAIM). But RAIM is designed to protect against faulted satellites, not spoofing attacks and only covers faults on a single satellite. The same limitation applies to the estimation scheme presented in [7]. Advanced RAIM (ARAIM) will go further by considering multiple faults and constellations [8].

Several techniques leverage spatial processing techniques to suppress spoofing signals when they are broadcasted from one direction. The malicious signals are removed from the received RF pattern, such that once again only authentic signals are visible to the receiver. This is achieved by steering a spatial null towards the largest power source using an array of multiple antennas [9], [10]. In the case of a moving receiver, a synthetic array can alternatively be constructed [11] and the correlation in Doppler variation leveraged to identify and exclude satellite signals coming from the same direction [12]. Several signal processing techniques have been explored that eliminate signals without the use of multiple or moving antennas. [13–15] eliminate one signal per PRN for example by a projection the signals onto their nullspace or by superpositioning the opposite signal. All three approaches however make strong assumptions on which signal per PRN is the authentic and which one is the spoofed signal.

Building upon our work in [16], in this paper we present techniques that offer continuity and integrity in the presence of a spoofing attack. Using a Software Defined Ratio (SDR), we track and decode both the authentic and spoofed signal for each PRN during an attack. Using the pseudorange residuals, we efficiently identify the two consistent (the authentic and spoofed) navigation solutions among all possible signal combinations. The question of which of the two consistent solutions to trust is then cast in the position domain, breaking with the main assumptions made in the literature.

For continuous navigation we suggest the use of a Multi Hypothesis Extended Kalman Filter (MHEKF). It coasts on an Inertial Measurement Unit (IMU) during the lift-off phase of an attack and identifies the authentic GNSS solution once two consistent position solutions are available. We show application examples using the TEXBAT dataset (described in detail in [17, 18]) and simulated attacks on driving data collected around Calgary.

To provide integrity during an ongoing spoofing attach we modify the Multi-Hypothesis Solution Separation (MHSS) algorithm developed for RAIM [19] to generate protection levels that encompass both the authentic and spoofed position solution while still protecting against satellite failures. We offer a larger, conservative bound as well as a tighter bound once the two consistent solutions have been identified. We once again show examples applying the bound to simulated attacks on driving data.

The remainder of the paper is organized in three main sections plus a summary and conclusion. In Section II we review the work in [16] on tracking multiple signals for each PRN and identifying consistent navigation solutions. In Section III we present the MHEKF architecture to provide continuous navigation during the attack. In Section IV we then cast protection levels around the authentic and spoofed position solution to provide integrity.

II. SIMULTANEOUS PROCESSING OF SPOOFED AND AUTHENTIC SIGNALS

The main prerequisite for this paper's work is that the victim needs to be able to receive the authentic signals with sufficient strength such that they can be tracked and decoded by the receiver. This is not necessarily the case during jam-then-spoof attacks, very high power advantage lift-off attacks or in the case of a nulling attack or physical signal blockage [3]. Once detected, the effect of a spoofing attack can in these cases only be reduced to a denial of service.
We further assume that a spoofing attack can be detected robustly before any error is introduced in the navigation solution, e.g. using a combination of metrics as detailed in [5].

Should the authentic signals be visible during the attack, the victim receives at least the authentic and the spoofed signal (and possible reflections thereof) for each spoofed PRN in view. The presence of more than one peak in the complex ambiguity function has been used for spoofing detection [20, 21]. Several approaches have been presented to track both signals simultaneously. In [22] and further developed in [23] through the use of auxiliary tracking channels, which is very similar to our strategy in [16].

In this section we briefly review the acquisition of multiple signals per PRN detailed in [16] and summarize the identification of the two consistent sets of signals, representing the spoofed and authentic navigation solution.

1. Multiple Signals per PRN

The cited techniques are generally capable of tracking multiple signals for each PRN if their code phase differs by at least one chip. For GPS L1 signals this corresponds to a pseudorange difference of around 300 m. If two signals are spaced more closely, the two correlation peaks overlap and no two distinct peaks are visible. It is worth noting that for GPS signals on the L5 frequency this overlap is reduced to around 30 m due to the 10x higher chipping rate of the signals. In Figure 1 (reproduced here from [16]) we can observe the in-phase and quadrature correlator tap values as recorded by an SDR during TEXBAT scenario 3. TEXBAT scenario 3 is a lift-off attack, during which the spoofer initially transmits signals with the same code delay τ and thereby range information as the authentic signal, overpowers the authentic signals, and then gradually changes τ. Before and during the early stage of the attack, the receiver's correlation peak shows no distortion as depicted in Figure 1a. Once the attacker changes τ, the correlation peak gets distorted by the superposition of correlation values (Figure 1b). This distortion is visible both in the inphase correlation due to the altered τ as well as the quadrature correlation due to the difference in Doppler and carrier phase between authentic and spoofed signal. Once the signals are spaced more than approximately one chip, a secondary peak is detected and tracked with a second set of 21 correlator pairs. This can be seen in Figure 1c.

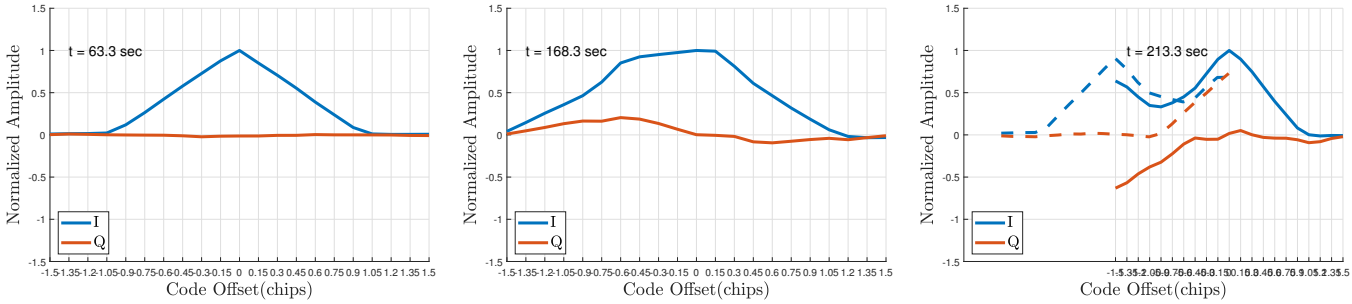

(a) Early phase, spoofer transmits replica of the authentic signal. No distortion visible.

(b) Lift-off phase, spoofer alters his τ and lures the victim off course. Distortion is visible in both I and Q channel.

(c) After lift-off. Two distinct but distorted peaks are visible. The secondary peak is indicated by dashed lines.

Figure 1: Correlator values during three stages of a lift-off spoofing attack. Plots based on 21 correlator pairs spaced evenly between ± 1.5 chips around the prompt correlator of identified peaks. All values are normalized by the prompt inphase correlator of the main peak.

The capability of tracking multiple signals especially for large power differences between the signals depends on several factors, such as the used front end and the receiver's analog to digital conversion bit resolution. The used SDR has a 16 bit resolution and was capable of tracking secondary signals during the 10 dB power advantage attack of TEXBAT scenario 2.

Once two peaks are detected, each is processed separately for an independent set of code phase, carrier phase and Doppler measurements as well as independent navigation data.

2. Two Consistent Solutions

Decoding multiple signals for each PRN offers a lot of opportunities, but comes with a caveat: the number of possible navigation solutions skyrockets. For n satellites in view with for example two signals each, we can compute 2^n solutions. Even if we leverage similarities through rank one updates, this quickly becomes impractical to be done in real time in a receiver.

Among the large number of possible navigation solutions, only few are actually of interest: the authentic and spoofed solutions. From here on out we limit our considerations to at most two signals per PRN, representing one authentic and one spoofed signal. Among the 2^n possible solutions, only 2 will be formed from a consistent set of pseudoranges.
Here we break with a common assumption in the literature: we do not assume anything about the spoofer's signal being stronger than the authentic. Either one of the two signals decoded for a specific PRN could be trustworthy or not. If one can be marked as compromised due to e.g. unrealistically high signal strength or disagreeing code and carrier rates it makes the following considerations easier, but it is not required.

a) A greedy Histogram Filter

Instead of computing 2^n solutions at every epoch, we follow a greedy approach to identify consistent sets of pseudoranges. In [16] we detail the approach using a histogram filter [24]. The filter keeps track of the likelihood of each considered signal combination at every epoch through a transition step and a recursive belief update.
We start at epoch 1 with any chosen signal combination c, e.g. by computing a navigation solution based on the strongest signals for each PRN and one based on all the weakest signals. Since only one signal combination is considered, its probability is $p(c) = 1$. The transition step then assigns probabilities to all "adjacent" signal combinations of $p(c)\lambda^\delta$. λ is a transition probability hyperparameter, in our experience 0.01 has worked well. δ is the number of signal assignments that are different between the new set and the original signal combination. We consider the "adjacent" combinations that differ only in $\delta = 1$ satellite assignment. This is the first step of the local greedy search of consistent signal combinations.
During the recursive belief update the likelihood of each considered combination is computed by forming the residual χ^2 statistic. Consider the linearized measurement equation in [25]

$$y = Gx + \varepsilon \tag{1}$$

where $y \in \mathbb{R}^n$ are the pseudorange measurements minus the expected range, $G \in \mathbb{R}^{n \times p}$ is the geometry matrix for p states ($p = 4$ for one constellation) in the state vector x. The measurement noise is normally distributed with covariance matrix $W^{-1} \in \mathbb{R}^{n \times n}$

$$W = \begin{bmatrix} \sigma_1^2 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \sigma_N^2 \end{bmatrix} \tag{2}$$

the χ^2-statistic t is given by [26]

$$t = y^T(W - WG(G^TWG)^{-1}G^TW)y \tag{3}$$

We can compute the probability of a vector of measurements for a combination c by evaluating the probability density function (pdf) of the central χ^2 distribution with $n - p$ degrees of freedom at t

$$p(y|c) = \frac{1}{2^{(n-p)/2}\Gamma((n-p)/2)} t^{(n-p)/2-1} e^{-t/2} \tag{4}$$

where Γ is the gamma distribution and t is computed using Eq. (3). Using Bayes Rule we can then compute the posterior probability of a signal combination c being consistent.

$$p(c|y) = \frac{p(y|c)p(c)}{\sum_{c \in C} p(y|c)p(c)} \tag{5}$$

where C is the set of all considered combinations (where the number of considered combinations $|C| \ll 2^n$) and $p(c)$ is the probability of the combination after the transition step.

If we continued these two steps at each epoch, $|C|$ would quickly grow until it reaches 2^n. To avoid this, we prune the number of considered combinations after every epoch. Possible pruning approaches are to consider only combinations with a posterior probability above some threshold, or to consider only the M most probable combinations. We choose the latter approach and limit $M = 20$ as this directly controls the computational effort necessary at each epoch.

Several measures can be taken to further limit the computational burden at every epoch. If the broadcasted ephemeris is consistent among all decoded signals, the satellite orbits only need to be propagated once per epoch. The signal combinations further use similar sets of pseudoranges; this should be leveraged through the use of rank one updates (as explained for example in Appendix I of [8]).

b) Application Examples

The result of this step are, at every epoch, a set of navigation solutions. Each has a probability of being computed from a consistent set of signals associated with it. Let us finish the brief but theoretical description of the approach with an application example. Specifically, we show the result when running the approach on TEXBAT scenarios 4 and 6 previously depicted in [16]. Both scenarios are position push attacks with 0.4-1 dB power advantage. During scenario 4 the receiver is stationary, during scenario 6 the receiver is moving. During both scenarios the attacker introduces a 600 m offset in the ECEF z coordinate. This is reflected in the navigation solutions depicted in Figure 2. The figures show the truth as a black line, and the solution obtained by an SDR without any spoofing defense in blue. Red dots show the obtained navigation solutions at every epoch, color-coded for each solution's posterior likelihood. During the initial, nominal period of the scenario only one solution exists. Once lift-off begins around 160 sec in Figure 2a and 130 sec in Figure 2b, multiple solutions begin to emerge, all more or less equally inconsistent as they were obtained from strongly distorted correlation peaks. Once lift-off is completed for all satellites around 260 sec (scenario 4, Fig. 2a) and 230 sec (scenario 6, Fig. 2b), the two consistent solutions of all nominal and all spoofed signals emerge. The dynamic scenario overall shows noisier conditions, likely because the number of visible satellites changes constantly, prompting the histogram filter to re-converge on consistent signal combinations.

III. CONTINUITY

In the previous section we have shown how to compute both the authentic and spoofed navigation solution among all possible signal combinations. Two problems persist however if we want to navigate continuously:

- For a significant period during lift-off we are without having identified any consistent navigation solution. We can not use satellite navigation.
- Once two consistent solutions have emerged, we still have not determined which one of the two is to be trusted.

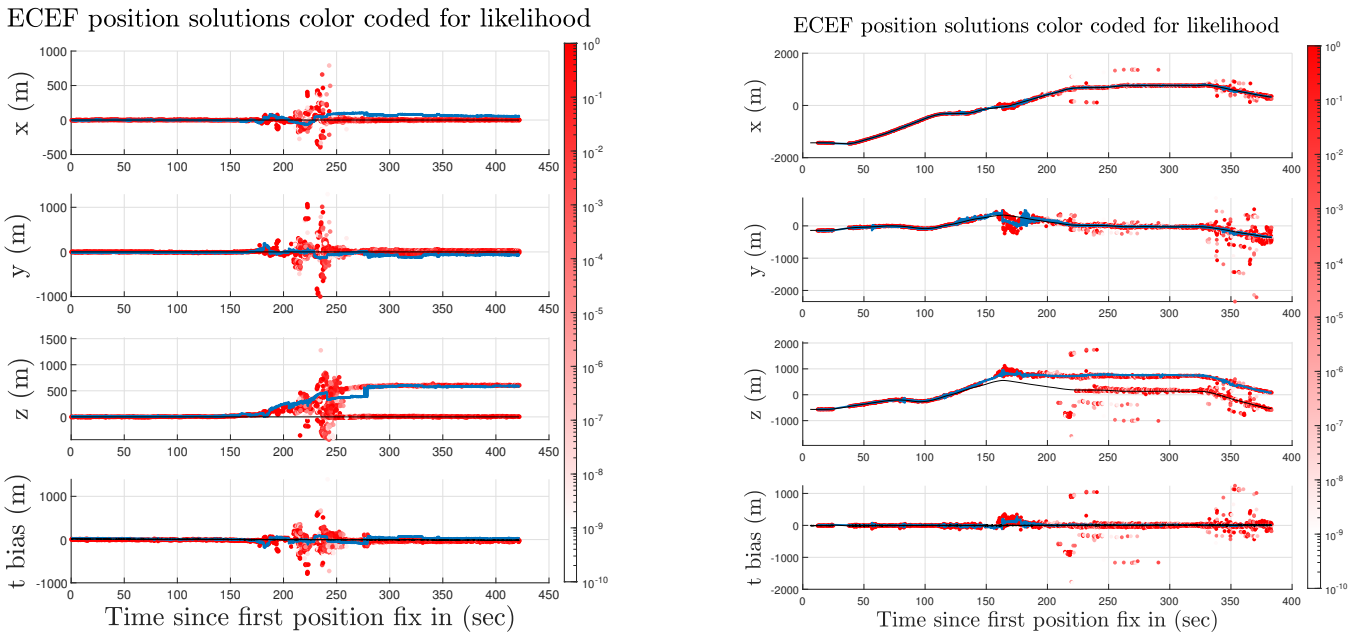

(a) Scenario $ds4$.

(b) Scenario $ds6$.

Figure 2: ECEF position solutions in meter during TEXBAT scenarios $ds4$ and $ds6$. Red dots show position solutions of the most probable solutions at each epoch, color coded for their likelihood as determined by the Histogram Filter architecture. The solid black line represents the truth. The blue solid line is the solution obtained if all main peaks are used, representing a receiver without spoofing defense.

1. The Multi-Hypothesis EKF

We resolve both issues using a Multi-Hypothesis Extended Kalman Filter (MHEKF) and an additional sensor allowing us to perform dead-reckoning navigation, such as an IMU. A similar approach is presented in [27], where the Doppler measurements are assumed unaffected and used during the coasting period. Here we use a loosely coupled EKF combining IMU and GNSS measurements presented in [16], based on the error-state implementation described in [28]. Each filter's state $x \in \mathbb{R}^{17\times 1}$ is given by

$$x = \begin{bmatrix} \Psi \\ r \\ \dot{r} \\ b_a \\ b_g \\ b_c \\ \dot{b}_c \end{bmatrix} \tag{6}$$

x contains the attitude in Euler angles $\Psi \in \mathbb{R}^{3\times 1}$, the ECEF position and velocity $r \in \mathbb{R}^{3\times 1}$ and $\dot{r} \in \mathbb{R}^{3\times 1}$, accelerometer and gyro bias terms $b_a \in \mathbb{R}^{3\times 1}$ and $b_g \in \mathbb{R}^{3\times 1}$ as well as the clock bias and clock drift terms b_c and $\dot{b}_c$. EKFs coupling IMU and GNSS measurements have been well studied and the interested reader is referred to the cited literature. In this section we will focus on the specifics of tracking multiple hypothesis with two filters.

During nominal conditions, one EKF tracks the antenna's state using both IMU and GNSS measurements. Standard system dynamics and IMU models are employed. The interested reader is referred to [16, 28], details are omitted here for brevity.

Once a spoofing attack is detected, the filter stops applying GNSS measurement updates and coasts on the IMU. Additional information from other sensors, e.g. such as velocity from an odometer or velocity and/or altitude from a pitot-static system could greatly improve the coasting capability but are not applied in this paper.

A second, "compromised" EKF in the meantime keeps employing GNSS measurement updates. Once the histogram filter described in Section II.2 a) has converged and identified the same two signal combinations as most probable in multiple consecutive epochs, the "trusted" EKF can resume using satellite navigation. To determine which GNSS solution to use, the two EKFs perform a localization with unknown correspondences [24]. Several approaches exist to this standard problem, here we employ a maximum likelihood estimator to assign the correspondences. The assignment $j(i)$ of the ith GNSS solution is

chosen to minimize a Mahalanobis distance.

$$j(i) = \arg\min_k (z_G^{(i)} - C_G x^{(k)})^T (C_G \Sigma^{(k)} C_G^T + R_G^{(i)})^{-1} (z_G^{(i)} - C_G x^{(k)}) \tag{7}$$

for state estimate of the kth EKF after the dynamics update $x^{(k)}$ and its covariance matrix $\Sigma^{(k)}$, the measurement of the ith GNSS solution $z_G^{(i)}$ and its covariance $R_G^{(i)}$ and the GNSS measurement matrix $C_G \in \mathbb{R}^{4\times 17}$ here given by

$$C_G = \begin{bmatrix} 0_{(3\times3)} & I_{(3)} & 0_{(3\times9)} & 0 & 0 \\ 0_{(1\times3)} & 0_{(1\times3)} & 0_{(1\times9)} & 1 & 0 \end{bmatrix} \tag{8}$$

where $0_{(a\times b)}$ is a matrix of zeros with dimensions $a \times b$ and $I_{(c)}$ is a $c \times c$ identity matrix.

The "compromised" EKF has been tracking the spoofed position solution during lift-off. At the moment that two consistent GNSS solutions become available, its estimate is likely close to the spoofed solution with tight covariance. The "trusted" EKF's state estimate on the other hand is likely closer to the authentic GNSS solution but with a large covariance. The precision of its estimate and size of its covariance matrix depend on the duration of the coasting phase, as well as the number and quality of the available non-GNSS sensors and the dynamics model.
In the case of a large error covariance and unprecise estimate, the confident estimate of the "compromised" EKF close to the spoofed position greatly helps choosing the correct assignment as it will strongly favor the spoofed solution for itself.

An exotic attack scenario could be imagined that takes advantage of this behavior. Consider a spoofer that broadcasts signals significantly *weaker* than the authentic signals, but strong enough to trigger the spoofing detection. During lift-off the receiver presumably keeps tracking the stronger, authentic signals. The "compromised" EKF tracks these authentic signals and claims them for itself during the correspondence assignment in Eq. (7). The "trusted" EKF with significantly larger covariance $C_G \Sigma^{(k)} C_G^T$ is then forced to use the spoofed solution. The spoofing attack, even though it failed to capture the receiver's tracking correlators (or rather: *because* it failed to) would be successful; our initial assumption of a (mostly) successful lift-off attack was violated.
If the receiver designer chooses to provide robustness against this type of attack, the "compromised" EKF should be omitted. The "trusted" EKF, after the coasting phase, then has no competition for the authentic GNSS solution close to its own estimate. This approach does require higher coasting performance through a higher quality IMU or additional sensors, to avoid the EKF drifting too far to be able to decide between the two consistent GNSS solutions once they are available.

2. Application Examples

Let us now support the theoretical derivations with application examples. Once again we turn to the TEXBAT dataset. Unfortunately, TEXBAT comes without ground truth pvt information. Simulating IMU measurements for the dynamic scenarios 5 and 6 that are in agreement with nominal GNSS measurements is therefore difficult. Instead, we limit the TEXBAT analysis to the stationary scenarios 3 and 4. They offer realistic spoofing attacks but only simulated IMU measurements. Additionally, we analyze IMU+GNSS data collected on highways around Calgary, Canada and inject simulated spoofing signals in post processing as an example with highly realistic coasting behavior.

a) TEXBAT

As we have seen in Figure 2a, the lift-off attack in TEXBAT scenario 4 is initially not 100% successful; the SDR does not immediately track the spoofed position solution. This effect is even more drastic in scenario 3, the SDR never tracks the spoofer's solution. Scenario 3 is time-push attack once again with a low power advantage of around 1 dB introducing a 600 m bias in the time solution.
The used IMU model is detailed in [16] following the standard text [28]. The noise characteristics reflect a tactical grade IMU and are specified by the power spectral density (PSD). Table 1 summarizes the root PSDs of the accelerometer noise $\sqrt{S_a}$, gyro noise $\sqrt{S_g}$, accelerometer bias $\sqrt{S_{b_a}}$ and gyro bias $\sqrt{S_{b_g}}$ from [16].

Table 1: Tactical grade IMU model characteristics. S_a, S_g, S_{b_a} and S_{b_g} represent the PSDs of accelerometer and gyro noise and biases.

$\sqrt{S_a}$	$\sqrt{S_g}$	$\sqrt{S_{b_a}}$	$\sqrt{S_{b_g}}$
$100\mu g/\sqrt{Hz}$	$0.1^\circ/\sqrt{h}$	$100\mu g/\sqrt{Hz}$	$2 * 10^{-6} rad\ s^{-0.5}$

We show results for both scenarios in Figure 3. Specifically we show the change in ECEF position and clock bias since the start of the scenario. We once again depict the SDR solution in blue and the nominal truth in black. The "compromised" EKF

estimate is shown in red, the "trusted" EKF in green. Dashed lines indicate the $\pm 2\sigma$ uncertainty of the filter. The two figures contain a lot of valuable information.

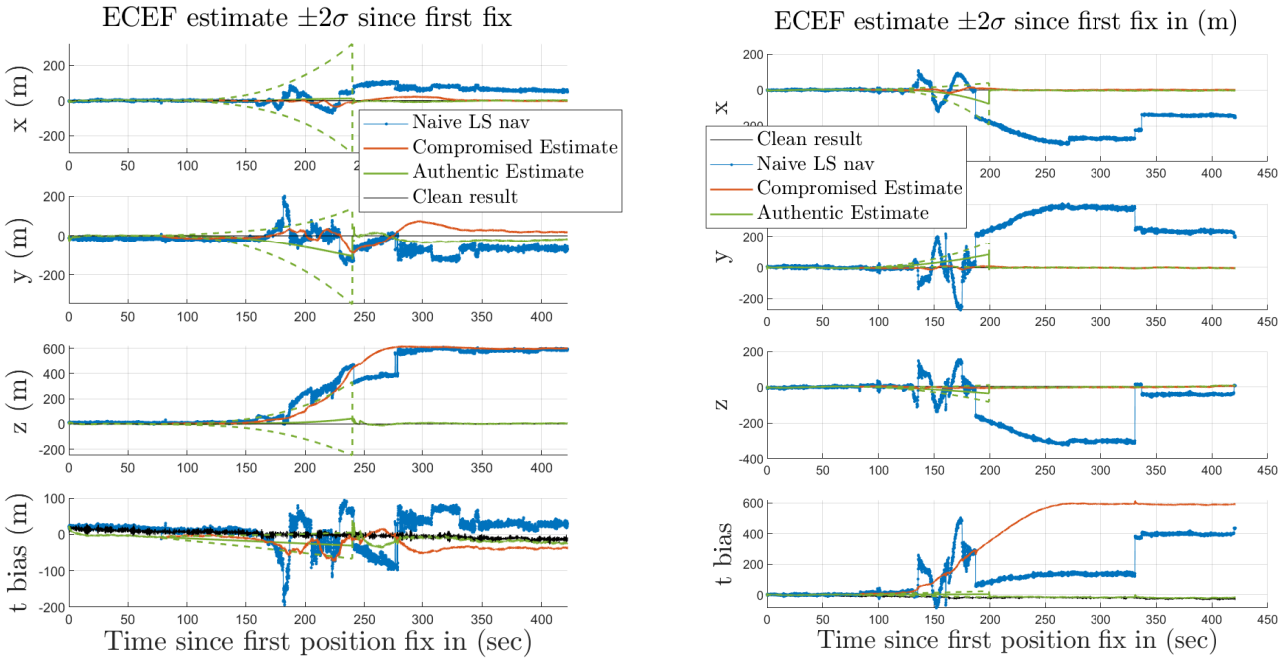

(a) Scenario 4.

(b) Scenario 3.

Figure 3: Navigation solution from clean data (black, our "truth"), a least squares solution using all main peaks (blue), an EKF tracking continuously (red) and an EKF coasting during times that a spoofing alarm is active (green). Dashed lines represent the $\pm 2\sigma$ bounds on the EKF estimates. The continuously tracking solution is compromised, it follows the spoofed signals. We can see the increase in position and clock estimate uncertainty during the coasting phase of the "authentic" solution which successfully tracks the true solution.

Most importantly, the green "trusted" solution closely tracks the authentic, black solution. We can see its uncertainty grow rapidly during the coasting phase. The coasting phase is around 40 sec longer in scenario 4, resulting in a significantly larger uncertainty.
The "compromised" EKF tracks the spoofed solution very well, closer than the SDR. A closer examination would reveal severely elevated IMU bias estimates during the lift-off phase, where the compromised GNSS measurements are in disagreement with the IMU measurements.

The results from Figure 3 underline the reliance of the approach on sufficient navigation capabilities during the coasting phase to support the assignment decision once two consistent GNSS solutions have emerged. This is largely dependent on how fast the lift-off is performed, as the uncertainty of the trusted measurement scales with time cubed.
As we stated in Section II.1, newer GNSS signals with higher chipping rate such as the GPS L5 signal should significantly help in this situation as secondary signals for each PRN appear 10 times faster.

b) Driving Data

The considered dataset was recorded from an automobile driving for one hour on highways around Calgary, Alberta, Canada. It has been reported on e.g. in the context of integrity for precise point positioning (PPP) solutions in [29]. Figure 4 shows the ground track. Truth pvt data is obtained from a NovAtel OEM729 paired with a tactical grade IMU as specified in Table 1 with forward and reverse processing.

The spoofing signals are injected 2100 seconds into the scenario. The IMU bias estimates in the EKF have converged at this point, providing good coasting performance. The attack scenario considers a highly potent spoofer with precise knowledge of the victim's planned trajectory, position and velocity. The spoofer performs a slow lift-off attack that moves along the same trajectory as the truth, but with a velocity increasing by $1\frac{mph}{s}$ for 10 sec until it reaches a velocity 10 mph faster than the truth. This slow lift-off results in a long coasting phase of around 240 sec until two signals are continuously decoded for at least 5 PRNs at 2340 sec. 3400 seconds into the scenario the attack stops, no more spoofed signals are broadcasted.

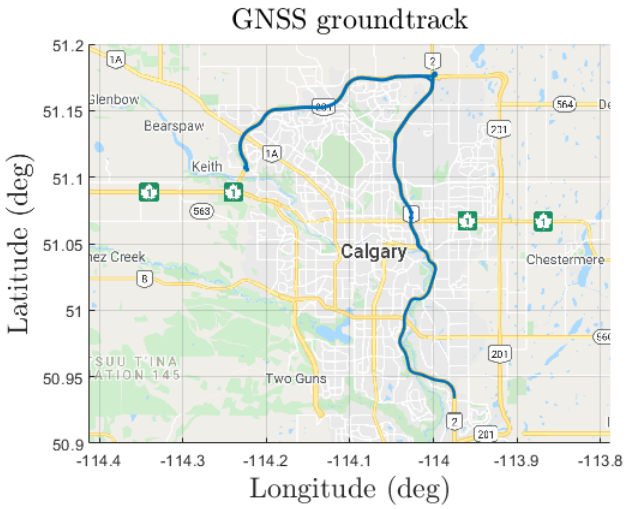

Figure 4: Ground track of the analyzed driving data.

Both EKFs use only GPS L1 signals at 1 Hz and IMU measurements at 100 Hz. The trusted EKF repeatedly applies a motion constraint (no slip) as pseudo-measurement [28] during the coasting phase.

We show results of both filter's estimates in Figure 5. In Figure 5a we can see the two EKF closely tracking both the true and spoofed positions during the attack. The "trusted", robust EKF navigates successfully despite the spoofing attack, whereas a standard, "compromised" EKF follows the spoofer's trajectory. During nominal conditions both filters track the true solution well.
In Figure 5b we take a closer look at the "trusted" EKF's performance. Specifically we depict its absolute navigation error and the 2σ bound given by its covariance in a local north (N), east (E), down (D) coordinate system. The error is generally very well bound by the 2σ bound, only a small bias remains for a while in the local vertical direction after two consistent solutions have been identified and satellite navigation has resumed at 2340 sec. We can see a significant reduction in the "trusted" filter's uncertainty and error at around 2250 sec. Here for a brief period a second signal was received for 5 satellites, allowing the trusted filter to use several GNSS measurements.

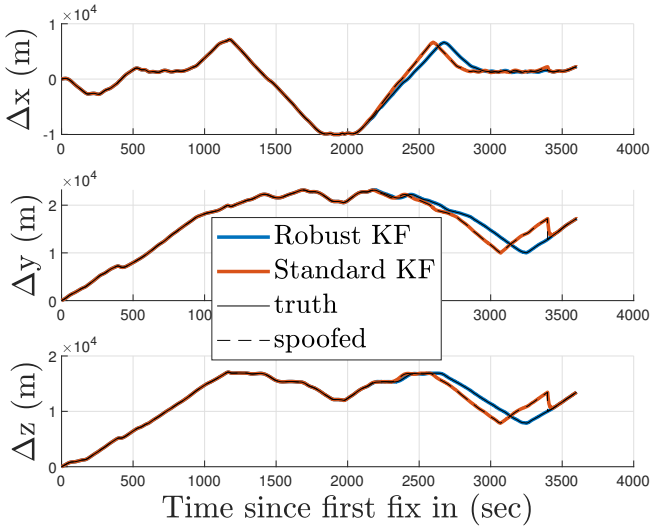

(a) Position truths and estimates of both EKFs relative to first epoch in ECEF coordinates.

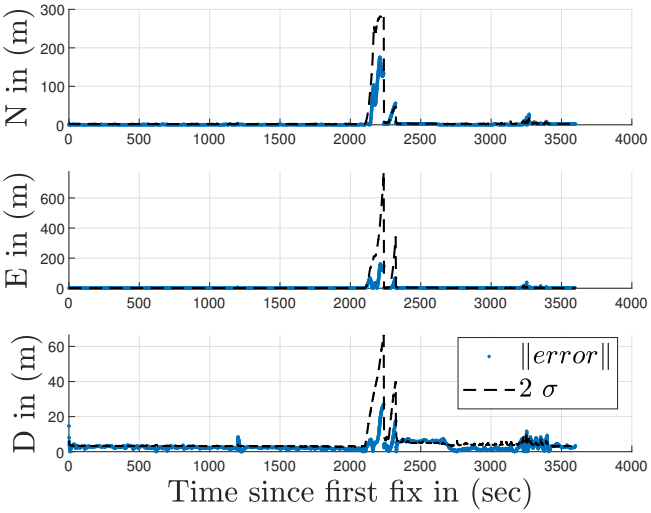

(b) Error and 2σ uncertainty of "trusted" EKF in local north (N), east (E), down (D) coordinates.

Figure 5: EKF results from the driving data scenario.

The results of the "trusted" EKF in Figure 5b represent continuous navigation despite the presence of a spoofing attack. The

temporarily significant navigation error and uncertainty can be further reduced through the use of signals with higher chipping rate such as GPS L5 or additional sensors for improved dead reckoning capabilities.

IV. INTEGRITY

The algorithms and methods we have explored in this paper offer continuous navigation thanks to identification of the spoofed and authentic navigation signals with low computational complexity. We have however employed some heuristics, and the 2σ bound in Figure 5b should not be interpreted as an integrity bound.
In this section we will now derive real time integrity bounds in the sense of RAIM, guaranteeing a certain low probability of hazardous misleading information (HMI, and its probability $PHMI$). We will see that these bounds can be derived directly from the Multi-Hypothesis Solution Separation (MHSS) equations in [19]. We then apply the algorithms to the driving data analyzed in Section III.2 b).

1. Protection Level Computation

We work with the definition of HMI used in e.g. [19]: it is the true position lying outside the error bound determined by the user. Its probability $PHMI$ can be computed as the sum of the integrity risk of all N considered threat or fault hypothesis H_i with $i = 0, \ldots, N$ and H_0 being the fault-free hypothesis weighted by their prior probability $P_{ap,i}$:

$$PHMI = \sum_{i=0}^{N} P_{ap,i} P(HMI|H_i) \tag{9}$$

Closely following the explanation in [19], we compute a subinterval L_i for each fault hypothesis such that

$$P(||x^{(i)} - x|| \geq L_i|H_i) \leq \frac{P(HMI|H_i)}{P_{ap,i}} \tag{10}$$

Here $x^{(i)}$ is the position solution computed using the un-faulted measurements under hypothesis H_i and x is the true position. Depending on the application, Eq. (10) can be phrased for each coordinate axis separately, for horizontal and vertical bounds, or for a 3-dimensional bound. In this paper we follow the aviation-implementation of RAIM and compute separate bounds for the local north, east and down direction.

The real time protection level (PL) is then computed for each direction as

$$PL = \max_i \left(||x^{(i)} - x^{(0)}|| + L_i \right) \tag{11}$$

where $x^{(0)}$ is the position estimate under the nominal hypothesis.

[19] offers a solution for L_i that satisfies Eq. (10). Let S be a weighted least squares estimator [25] to solve the linearized measurement equation (1) and $S^{(i)}$ the estimator for a fault-free satellite combination under hypothesis H_i.

$$x^{(i)} = S^{(i)} y \tag{12a}$$

$$S^{(i)} = \left(G_i^T W G_i \right)^{-1} G_i^T W \tag{12b}$$

The navigation solution uncertainty in the kth coordinate axis is then given by

$$\sigma_{k,i}^2 = S_{k,\cdot}^{(i)} W^{-1} S_{k,\cdot}^{(i)} = \left(G_i^T W G_i \right)_{k,k}^{-1} \tag{13}$$

where the subscript $k, \cdot$ indicates the kth row of the matrix. In RAIM pseudorange errors are modeled as Normally distributed with covariance W^{-1} and maximum bias b. The maximum position state error in the kth direction under the ith hypothesis introduced by biases is bounded by [19]

$$B_{k,i} = \sum_{l=1}^{n} |S_{k,l}^{(i)}| b_l \tag{14}$$

L_i in the kth coordinate direction can then be calculated as

$$L_{k,i} = K_{HMI,i}\sigma_{k,i} + B_{k,i} \tag{15}$$

with

$$K_{HMI,i} = Q^{-1}\left(1 - \frac{PHMI_i}{2P_{ap,i}}\right) \tag{16}$$

where Q^{-1} is the inverse Normal cumulative distribution function and $PHMI_i$ is the HMI budget allocated to the ith fault mode. The budgets are allocated such that they sum up to the HMI requirement.

$$PHMI_{req} = \sum_{i=0}^{N} PHMI_i \tag{17}$$

This concludes the real time PL computation using MHSS which we have largely taken from [19]. We now explore two options how to modify these considerations to include spoofing attacks in the threat model.

a) Conservative Approach: Inflated biases

In Section II.1 we have briefly reviewed how to track and decode multiple signals per PRN, resulting in more than one pseudorange measurement from each satellite. Two distinct peaks in the auto-correlation function can be tracked if they are at least one chip length apart [16]. In the case of a spoofing attack, the maximum distance between the fault free pseudorange measurement and any other pseudorange measurement for that satellite corresponds to a possible bias in that measurement. Specifically, we can update the maximum pseudorange bias of the lth bias to include the multiple peaks.

$$b_l^* = 2b_l + \max_c |\Delta y_c| \tag{18}$$

where we use the pseudorange bias $|\Delta y_p|$ between the fault free measurement $y^{(0)}$ and the measurements of pseudorange combination $c\ y^{(c)}$. If a spoofing alarm was raised but only one peak is detected, the signals must be less than one code chip apart. The pseudorange bias is then set to the chip length δ.

$$\max_c |\Delta y_c| = \begin{cases} \max\limits_p |y^{(c)} - y^{(0)}| & \text{if more than one peak detected} \quad (19a) \\ \delta & \text{otherwise} \quad (19b) \end{cases}$$

For GPS L1 signals, $\delta = (2.9979 \cdot 10^8)/(1.023 \cdot 10^6) = 293.05$ m.

With the updated biases, Equations (11 - 16) can then be applied as before to compute protection levels. These protection levels now include in their threat model all spoofing attacks that can be detected before any pvt error is introduced and that allow tracking of the authentic and spoofed signal once they are more than one chip apart.

Using this simple modification has one caveat. The protection level computation in RAIM leverages the linear approximation of the observation equation (1). This approximation is very accurate even up to several hundreds of meters. With pseudorange biases exceeding several hundred meter however, this approximation error grows. Preliminary results for protection levels of several km show this error on the order of magnitude of several cm. Within the range of this error, Eq. (14) possibly no longer bounds the position error introduced by the biases. Further research is necessary to determine the impact of the approximation error and how it could be mitigated.

b) Tighter Bounds around Consistent Solutions

The conservative approach lets us include a wide range of spoofing attacks in the threat model with a rather simple modification of the pseudorange error model. Using this modified error model in Eq. (14) bounds all 2^n (in the case of two signals per PRN) possible combinations of pseudorange measurements. As we have discussed in Section II.2, most of these solutions actually stem from highly inconsistent combinations of pseudoranges and can be discarded as improbable navigation solutions. Their likelihood as given by Eq. (4) is significantly smaller than the allowed $PHMI$. The bound calculated using the updated pseudorange bias from Eq. (18) is therefore very conservative.

Here we propose an alternative approach. Following the spirit of MHSS of considering multiple hypotheses, we cast RAIM protection levels around each consistent solution. A global protection level PL^* in the kth coordinate direction can then be computed to encompass the individual PLs.

$$PL_k^* = \max_c \left(|x_k^{(c)} - x_k^{(0)}| + PL_k^{(c)} \right) \tag{20}$$

with $c = 0, \ldots, |C| - 1$ for $|C|$ considered consistent sets of pseudorange measurements. Any set of measurements can be chosen as the "fault free" set $c = 0$.

The individual protection levels $PL^{(c)}$ are calculated as usual using Eq. (11 - 17) with HMI budgets that sum up to the overall HMI requirement.

$$PHMI_{req} = \sum_{c=0}^{|C|-1} PHMI_{req}^{(c)} \tag{21}$$

The individual PLs are therefore slightly inflated compared to standard RAIM, as the HMI budget is shared among all consistent solutions.

This approach allows for tighter protection levels but comes with a major assumption. It assumes that the set of all authentic satellite signals is among the considered sets. In the examples shown in the next section we make this assumption once the Histogram Filter described in Section II.2 a) has converged on the two most consistent solutions. The results of this approach are encouraging, but a proof for this assumption is left for future work.

2. Application to Driving Data

We now apply both protection level algorithms to the driving data analyzed in Section III.2 b). Once again we assume a spoofing detection mechanism in place that detects the attack before any bias is introduced in the pvt solution. All protection levels in this section are computed with $PHMI_{req} = 10^{-7}/hour$ and $P_{ap} = 10^{-5}$ following the guidance provided by [30].

We start the analysis with a single snapshot 2250 sec into the scenario. At this point, the lift-off phase of the attack is completed with two peaks visible in the autocorrelation function for several satellites. More than one consistent navigation solution can be computed by the receiver, and the histogram filter has converged on the two most consistent solutions. In Figure 6a we show the conservative and tight PLs computed using the approaches in Sections IV.1 a) and IV.1 b), respectively. The difference is quite dramatic, the tight protection levels are significantly smaller at this moment. Position solutions are once again depicted color-coded according to their likelihood and all lie within the tight PL.

Figure 6b shows a zoomed in view of the situation. We further depict the actual true and spoofed locations in the figure, as well as the PLs $PL^{(s)}$ around each individual consistent solution that make up the tight PL. Any solution with likelihood (given by Eq. (5)) $\geq PHMI_{req}$ is considered here.

The difference between conservative and tight PL is quite significant in the example shown in Figure 6. This is not always the case. The difference depends on various factors such as satellite geometry and the offset between authentic and spoofed solution. To analyze the difference between conservative and tight PL as well as the absolute navigation error in more detail than a single snapshot, we show the three parameters throughout the entire scenario in Figure 7. The navigation error here is the distance between the true position and the solution used in the "trusted" EKF described in Section III.1.

We can see both PLs easily bounding the error. Depending on the epoch, the tight PL is significantly smaller than the conservative bound, in other epochs the two are almost identical. A significant improvement is achieved in the vertical direction, because the spoofer barely introduces any vertical error. Based on these results, depending on the scenario, a user might very well elect to continue using satellite navigation despite an attack. The victim knows not only the likely difference between authentic and spoofed position solution, but has a guarantee on the maximum navigation error. Even a safety of life system such as an

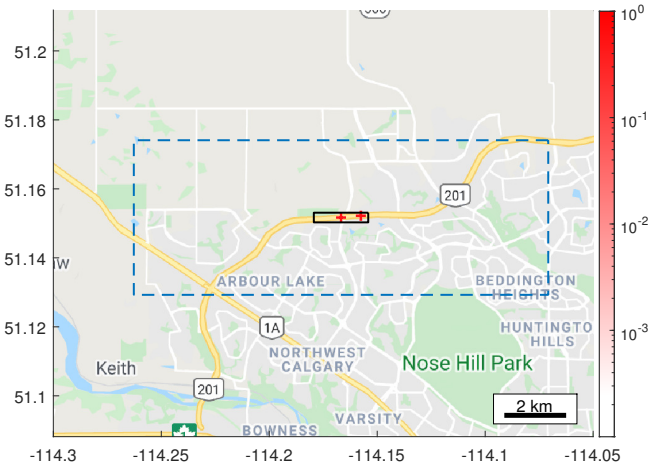

(a) Conservative (blue dashed line) and tight (black solid line) PLs together with all consistent position solutions.

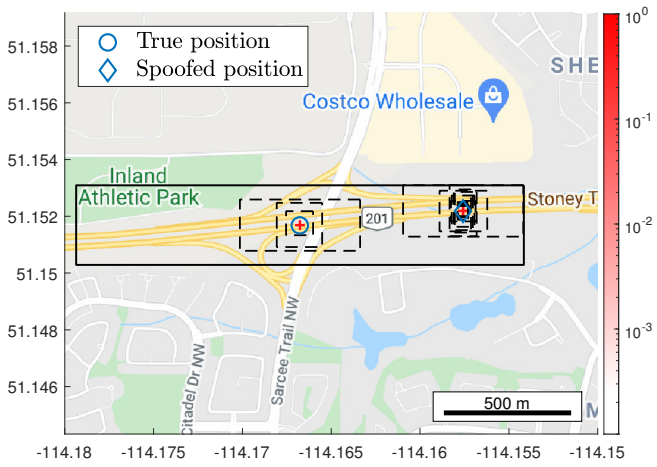

(b) Tight PL and individual PLs around consistent solutions. The true and spoofed location are indicated by a circle and diamond, with the diamond mostly hidden behind the individual PLs.

Figure 6: Protection levels shortly after lift-off. Position solutions are color-coded depending on their likelihood.

aircraft might still elect to use satellite navigation in a terminal area if the conservative PL is within the required navigational performance.

During the lift-off period between 2100 and around 2250 sec, no "trusted" GNSS solution or tight PL exists. The conservative PL however still provides an integrity bound on the spoofed solution using the pseudorange bias of one chip length δ as described in Section IV.1 a).

The navigation error shown in Figure 7 remains small throughout the scenario. We have leveraged an IMU in the MHEKF architecture to successfully identify the authentic navigation solution. This success is however not guaranteed, or an IMU might not be available. As a last result we therefore show the ratio of the worst case navigation error, if at every epoch the worst navigation solution (with the maximum offset) were chosen, and the tight PL in Figure 8. To bound the worst case errors, the ratio has to be < 1 at all times.

The figure shows successful bounding of the error, the integrity guarantee holds. The shown ratio does get remarkably close to 1 a significant number of times however. These are cases where the tight PL computation is dominated by the large distance between the individual solutions $\max_p |\Delta y_p|$. The individual $PL^{(s)}$ are comparably small, resulting in ratios close to 1.

V. SUMMARY AND CONCLUSION

In this paper we have shown techniques to provide both continuous navigation and protection levels during the presence of a spoofing attack. Continuity is achieved with the help of an IMU in a Multi-Hypothesis EKF, while integrity is provided by an extension to the Multi-Hypothesis Solution Separation algorithm developed in the scope of RAIM. The foundation for both approaches is the simultaneous reception of authentic and spoofed satellite signals. We demonstrate successful results when testing the algorithms against the TEXBAT dataset and a simulated attack on real IMU and GNSS data collected during a highway driving scenario.

Continuity, and especially integrity in the presence of a spoofing attack is a fairly unexplored research area and a lot of future work remains to be done. For example the effect of the linear approximation to the position solution over several km needs to be explored in more detail, just like a proof for the convergence of the greedy histogram filter. Further performance improvements would be possible when incorporating the MHEKF architecture into the integrity work. [29] is a good example for the integrity bounds possible when running multiple EKFs, one for each threat hypothesis. A tightly coupled EKF could be considered to extend the continuity and tight integrity bounds when two signals are received for less than 5 satellites.
It could further be interesting to explore the use of the techniques presented in this paper under multipath conditions, attempting to both identify the correct solution and provide an integrity bound on its error.

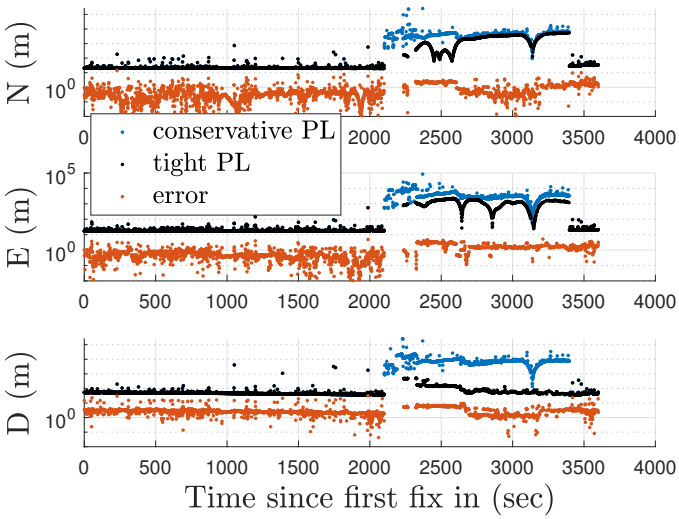

Figure 7: Conservative and tight PL as well as the navigation error in a local north (N), east (E), down (D) frame.

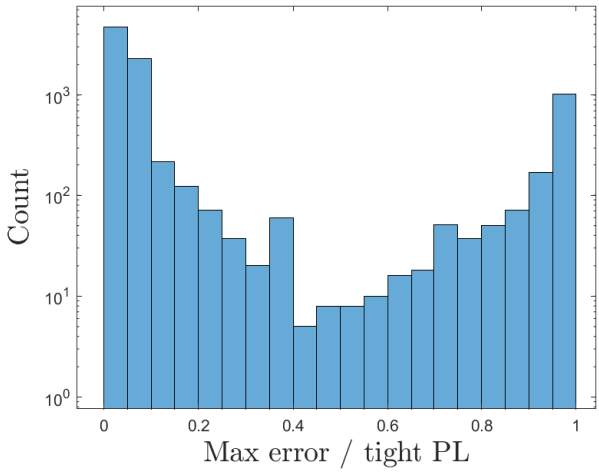

Figure 8: Ratio of worst case navigation error and tight PL.

ACKNOWLEDGMENTS

The authors thank the Federal Aviation Administration (FAA) and the Stanford Center for Position Navigation and Time (SCPNT) for sponsoring this research. We thank NovAtel for generously providing the driving data used in this paper.

REFERENCES

[1] A. Jafarnia-Jahromi, A. Broumandan, J. Nielsen, and G. Lachapelle, "GPS vulnerability to spoofing threats and a review of antispoofing techniques," *International Journal of Navigation and Observation*, 2012.

[2] C. Günther, "A Survey of Spoofing and Counter-Measures," *Navigation, Journal of the Institute of Navigation*, 2014.

[3] M. L. Psiaki and T. E. Humphreys, "GNSS Spoofing and Detection," *Proceedings of the IEEE*, vol. 104, no. 6, pp. 1258–1270, 2016.

[4] K. D. Wesson, J. N. Gross, T. E. Humphreys, and B. L. Evans, "GNSS Signal Authentication Via Power and Distortion Monitoring," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 2, pp. 739–754, 2017.

[5] F. Rothmaier, Y.-h. Chen, S. Lo, and T. Walter, "A Framework for GNSS Spoofing Detection through Combinations of

Metrics," *IEEE Transactions on Aerospace and Electronic Systems*, 2021.

[6] B. W. Parkinson and P. Axelrad, "Autonomous GPS Integrity Monitoring Using the Pseudorange Residual," *Navigation*, vol. 35, no. 2, pp. 255–274, 1988.

[7] Y. Oshman and M. Koifman, "Robust navigation using the global positioning system in the presence of spoofing," *Journal of Guidance, Control, and Dynamics*, vol. 29, no. 1, pp. 95–104, 2006.

[8] J. Blanch, T. Walter, P. Enge, Y. Lee, B. Pervan, M. Rippl, A. Spletter, and V. Kropp, "Baseline Advanced RAIM User Algorithm and Possible Improvements," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 1, pp. 713–732, 2014.

[9] S. Daneshmand, A. Jafarnia-Jahromi, A. Broumandan, and G. Lachapelle, "A low-complexity GPS anti-spoofing method using a multi-antenna array," in *25th International Technical Meeting of the Satellite Division of the Institute of Navigation 2012, ION GNSS 2012*, vol. 2, Nachville, TN, 2012, pp. 1233–1243.

[10] A. Konovaltsev, M. Cuntz, C. Haettich, and M. Meurer, "Autonomous Spoofing Detection and Mitigation in a GNSS Receiver with an Adaptive Antenna Array," in *26th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS+ 2013*. Nashville, TN: The Institute of Navigation, 2013, pp. 2937–2948.

[11] J. Nielsen, A. Broumandan, and G. Lachapelle, "GNSS spoofing detection for single antenna handheld receivers," *Navigation, Journal of the Institute of Navigation*, vol. 58, no. 4, pp. 335–344, 2011.

[12] A. Broumandan, A. Jafarnia-Jahromi, and G. Lachapelle, "Spoofing detection, classification and cancelation (SDCC) receiver architecture for a moving GNSS receiver," *GPS Solutions*, vol. 19, no. 3, pp. 475–487, 2015.

[13] S. Han, L. Chen, W. Meng, and C. Li, "Improve the Security of GNSS Receivers Through Spoofing Mitigation," *IEEE Access*, vol. 5, pp. 21 057–21 069, 2017.

[14] F. Wang, H. Li, and M. Lu, "GNSS Spoofing Detection and Mitigation Based on Maximum Likelihood Estimation," *Sensors (Switzerland)*, vol. 17, no. 7, 2017.

[15] Y. Guo, L. Miao, and X. Zhang, "Spoofing Detection and Mitigation in a Multi-correlator GPS Receiver Based on the Maximum Likelihood Principle," *Sensors (Switzerland)*, vol. 19, no. 1, p. 17, 2019.

[16] F. Rothmaier, Y.-h. Chen, S. Lo, and T. Walter, "GNSS Spoofing Mitigation in the Position Domain," in *Proceedings of the 2021 International Technical Meeting of The Institute of Navigation*, 2021, pp. 42 – 55.

[17] T. E. Humphreys, J. A. Bhatti, D. Shepard, and K. D. Wesson, "The Texas spoofing test battery: Toward a standard for evaluating GPS signal authentication techniques," in *Proceedings of the 25th International Technical Meeting of the Satellite Division of the Institute of Navigation 2012, ION GNSS 2012*, Nashville, TN, 2012, pp. 3569–3583.

[18] C. A. Lemmenes, P. Corbell, and S. Gunawardena, "Detailed analysis of the TEXBAT datasets using a high fidelity software GPS receiver," in *29th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS 2016*, vol. 5, 2016, pp. 3027–3032.

[19] J. Blanch, T. Walter, and P. Enge, "RAIM with Optimal Integrity and Continuity Allocations Under Multiple Failures," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 46, no. 3, pp. 1235–1247, 2010.

[20] L. Scott, "Anti-Spoofing and Authenticated Signal Architectures for Civil Navigation Systems," *Proceedings of the 16th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS/GNSS 2003)*, 2003.

[21] C. Hegarty, B. W. O'Hanlon, A. Odeh, K. Shallberg, and J. Flake, "Spoofing detection in GNSS receivers through cross-ambiguity function monitoring," in *Proceedings of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS+ 2019*. Miami, Florida: Institute of Navigation, 2019, pp. 920–942.

[22] A. Ranganathan, H. Ólafsdóttir, and S. Capkun, "SPREE: A spoofing resistant GPS receiver," in *Proceedings of the Annual International Conference on Mobile Computing and Networking, MOBICOM*, 2016, pp. 348–360.

[23] C. Yang and A. Soloviev, "All Signal Acquisition Processing for Spoofing Detection , Estimation , Mitigation and Intent Analysis," in *Proceedings of the 33rd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2020)*. Institute of Navigation, 2020, pp. 3338 – 3351.

[24] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. Cambridge, Massachusetts: The MIT Press, 2005.

[25] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance; Revised Second Edition*, 2nd ed. Lincoln, Massachusetts: Ganga-Jamuna Press, 2011.

[26] M. Joerger, F. C. Chan, and B. Pervan, "Solution separation versus residual-based RAIM," *Navigation, Journal of the Institute of Navigation*, vol. 61, no. 4, pp. 273–291, 2014.

[27] M. Berardo, E. G. Manfredini, F. Dovis, and L. Lo Presti, "A spoofing mitigation technique for dynamic applications," in *2016 8th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing, NAVITEC 2016*. IEEE, 2017, pp. 1–7.

[28] P. D. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*, 1st ed. Boston, NY: Artech House, 2008.

[29] K. Gunning, J. Blanch, T. Walter, L. De Groot, and L. Norman, "Integrity for Tightly Coupled PPP and IMU," in *Proceedings of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS+ 2019*, 2019, pp. 3066–3078.

[30] Working Group C, "EU-US Cooperation on Satellite Navigation Working Group C, ARAIM Technical Subgroup, Milestone 3 Report," Tech. Rep. 1, 2016. [Online]. Available: https://www.gps.gov/policy/cooperation/europe/2016/working-group-c/ARAIM-milestone-3-report.pdf