Received: 27 October 2021 | Revised: 2 June 2023 | Accepted: 19 June 2023

DOI: 10.33012/navi.616

ORIGINAL ARTICLE

Stochastic Reachability-Based GPS Spoofing Detection with Chimera Signal Enhancement

Tara Mina | **Ashwin Kanhere** | **Shreyas Kousik** | **Grace Gao**

Department of Aeronautics and Astronautics, Stanford University

Correspondence
Grace Gao
Department of Aeronautics and Astronautics, Stanford University
Stanford, CA, USA, 94305
Email: gracegao@stanford.edu

Abstract

To protect civilian global positioning system (GPS) users from spoofing attacks, the U.S. Air Force Research Lab has proposed the chips-message robust authentication (Chimera) enhancement for the L1C signal. In particular, the Chimera fast channel allows users to authenticate the received GPS signal once every 1.5 or 6 s, depending on the out-of-band source utilized for receiving the fast channel marker keys. However, for many moving receiver applications, receivers often use much higher GPS measurement rates, at 5–20 Hz.

In this work, we derive a stochastic reachability (SR)-based detector to perform continuous GPS signal verification and state estimation between Chimera authentications. Our SR detector validates the received GPS measurement against any self-contained sensor, such as an inertial measurement unit, in the presence of bounded biases in the sensor error distributions. We demonstrate via Monte Carlo simulations that our detector satisfies a user-defined false alarm requirement during nominal conditions, while successfully detecting a simulated spoofing attack. We further demonstrate that our SR state estimation filter successfully bounds the true state during both authentic and spoofed conditions.

Keywords

Chimera, formal verification, Kalman filter, probabilistic zonotopes, spoofing detection, stochastic reachability

1 | INTRODUCTION

To provide secure navigation for civilian global positioning system (GPS) users, the Air Force Research Lab (AFRL) has developed the chips-message robust authentication (Chimera) (Anderson et al., 2017) signal enhancement for the GPS L1C signal (GPS Directorate, 2022). Chimera inserts a digital signature within both the navigation message and the pilot channels of L1C to allow civilian users to jointly authenticate both components of the signal (AFRL Space Vehicles Directorate, Advanced GPS Technology, 2019). The AFRL will broadcast and test this signal enhancement on the upcoming Navigation Technology Satellite 3 experimental platform, which will be launched in 2024 (Cozzens, 2021; Divis, 2019, AFRL, 2023). If incorporated within the GPS L1C signal, the Chimera enhancement will be the first GPS signal encryption scheme available for civilian users, thereby enabling secure navigation for all future GPS users.

To ensure that the GPS signal cannot be forged by a malicious attacker, the Chimera-enhanced satellite segment will only publish the encryption key to the user segment after the subsequent key has already been updated. Users with access to only the GPS L1C signal receive the slow channel encryption key once every 3 min within the GPS L1C navigation message. However, users with access to secure out-of-band channels will be able to receive the fast channel encryption key once every 1.5 s, e.g., through a secure internet connection, or once every 6 s, e.g., through an augmentation system (Cozzens, 2021; GPS Directorate, 2022). With these encryption keys, users can authenticate their received GPS signal periodically at the rate of key reception. However, in either case, the Chimera signal authentication feature is not continuously available. In particular, even fast channel users will experience a 6-s latency in signal authentication, whereas GPS position update rates for moving receivers, such as autonomous vehicles, are typically 5–20 Hz.

To address this challenge, the present work develops a method to provide continuously available authenticated navigation solutions using Chimera. In particular, we focus on using the Chimera authentication feature in this work, because of the availability of its detailed interface specification (AFRL Space Vehicles Directorate, Advanced GPS Technology, 2019); however, the techniques and derivations in this work can be applied to any setting in which periodic authentication information is available. In addition to the Chimera authentication information, we utilize measurements from another self-contained sensor on-board the vehicle, such as an inertial measurement unit (IMU), in order to validate the received GPS signal, while accounting for measurement uncertainties and unknown bounded biases in the self-contained sensor and GPS measurements under authentic conditions. This work is based on our recent ION GNSS+ 2021 conference paper (Mina et al., 2021).

1.1 | Related Work

Prior work has been conducted to perform signal verification on encrypted global navigation satellite system (GNSS) signals to detect the reception of a false received signature as well as to detect a secure code estimate-and-replay (SCER) attack, where an attacker rapidly estimates the encrypted spreading signals and immediately rebroadcasts them to mimic authentic GNSS signals on the fly (Humphreys, 2013; Wesson et al., 2012). Indeed, these detection strategies are of critical importance to leverage the security benefits of cryptographically secured GNSS signals as well as to defend against more sophisticated SCER attacks. However, these techniques do not address simpler spoofing attacks during the critical authentication interval before the encrypted signature is received.

Additional prior work has utilized point-valued state estimation to conduct spoofing detection with on-board inertial navigation system measurements by monitoring Kalman filter innovations in a tightly coupled system (Tanıl et al., 2017) and by comparing IMU-estimated state trajectories with GPS-estimated trajectories (Broumandan & Lachapelle, 2018). However, these detectors often require strong assumptions regarding the underlying distributions of sensing uncertainty, such as known sensor biases and nominal unbiased Gaussian measurement noise.

Many approaches for conservative error modeling have been developed for the analysis of GNSS integrity systems. These approaches include cumulative distribution function (cdf) overbounding, which uses a single Gaussian distribution to bound both sides of the true error cdf (DeCleene, 2000), i.e., overbounding the left side, where the cdf is less than 0.5, and underbounding the right side, where the cdf is greater than 0.5. A desirable property of a conservative error modeling framework is the ability to overbound the summation of two independent error distributions using their individual overbounding distributions, i.e., via a convolution in the probabilistic domain. For GNSS integrity analysis, this property is particularly valuable for determining an overbounding distribution in the position domain, using individual overbounding distributions in the range domain (DeCleene, 2000; Rife et al., 2006; Rife et al., 2004). However, cdf overbounding only satisfies this summation bounding property for symmetric, zero-mean, and unimodal error distributions (DeCleene, 2000). To handle nonsymmetric and multimodal error distributions, paired overbounding techniques have been developed, in which two distributions are used to bound the left and right sides of the error cdf (Rife et al., 2006; Rife et al., 2004). Additional techniques have also been developed in which cdf overbounding and paired overbounding are combined to create an intermediate overbounding distribution and the cdf overbounding requirements are relaxed while maintaining an upper bound of the final position domain integrity risk (Blanch et al., 2017, 2018). However, all of these past works typically only model overbounding scalar distributions and assume knowledge about the exact underlying error distribution, such as the distribution of sensing biases (e.g., GNSS multipath errors), which can lead to difficulties in reliable modeling in practice and for all types of sensing errors.

Promising work has also been conducted in utilizing formal verification techniques, including stochastic reachability (SR), in the context of safe satellite-based navigation applications. SR, which is described in greater detail in Section 2, addresses the challenges of point-valued state modeling approaches by using less restrictive error models (e.g., unknown, but bounded, biases in sensing uncertainty). Indeed, unlike many prior works on conservative error modeling for GNSS integrity analysis (Blanch et al., 2018; DeCleene, 2000; Rife et al., 2006), SR provides a complete probabilistic overbound over a set of possible vehicle state distributions. Additionally, SR provides an elegant framework to extend this probabilistic overbounding to a multidimensional space in an efficient manner that provides tighter bounds than maintaining overbounds in each dimension separately, as discussed in Section 2. Furthermore, set representations for SR have been developed, which can be used to efficiently evaluate the overbound under linear mapping and summation operations (Althoff, 2010; Althoff et al., 2009), as discussed in more detail in Section 2. In prior works, this technique has been used to compute a set of reachable unmanned aerial vehicle states for safe trajectory planning, while incorporating bounds on potential GNSS measurement biases (Shetty & Gao, 2019, 2021). Additionally, formal verification techniques through SR have been utilized to perform secure GPS timing estimation within a network of phasor measurement unit devices in a power grid network (Bhamidipati & Gao, 2020a).

1.2 | Overview of Proposed Method and Contributions

We propose a spoofing detector to provide continuous GPS signal verification between Chimera authentication times using SR analysis, inspired by recent methods such as those described in Althoff et al. (2009) and Bhamidipati and Gao (2020a, 2020b). We derive our spoofing detector and state estimator for a generic linear or nonlinear self-contained sensor model with GPS positioning measurements. To experimentally validate our technique, we implement our algorithm for a ground receiver paired with (1) a linear sensor model of two-dimensional (2D) acceleration inputs in the navigation frame of reference as well as (2) an on-board IMU sensor as the self-contained sensor. For each time instant at which the receiver position is updated, our formal verification method leverages the previously authenticated set of Chimera measurements in combination with conservative error models for the GPS and self-contained sensor measurements to ensure that the detector meets a user-defined false alarm threshold on declaring a spoofing event.

To address the challenges of point-valued spoofing detection methods and to leverage the Chimera signal enhancement, our proposed formal verification technique:

1. enables continuous GPS signal verification between Chimera authentication times by validating the received signal against local self-contained sensors,
2. provides a probabilistic overbound on the set of possible vehicle states for navigation, in the presence of both stochastic uncertainties and bounded measurement biases for the self-contained sensor and GPS sensor during authentic conditions, and
3. evaluates a spoofing detection statistic that satisfies a user-defined false alarm metric, while accounting for potential biases in the self-contained sensor and GPS measurements during nominal unspoofed operation.

1.3 | Paper Organization

In Section 2, we establish our notation and key definitions. In Section 3, we present our proposed SR filter and spoofing detector. We present our experimental results in Section 4 and conclude in Section 5.

2 | PRELIMINARIES

In this section, we introduce the notation used throughout the paper. We also provide a basic introduction to probabilistic zonotopes (p-zonotopes), which are used for representing stochastic sets. Other mathematical objects have been developed, including ellipsotopes (Kousik et al., 2021), which can provide tighter bounding stochastic sets than p-zonotopes. However, the primary objective of this work is to develop the overall framework and analysis for an SR-based spoofing detector and state estimator. As a result, we focus on using a single stochastic set representation for our derivations; in particular, we utilize p-zonotopes. Zonotopic set representations are a popular choice for modeling reachable sets in formal verification applications (Kousik et al., 2019; Medina Lee et al., 2019; Schürmann et al., 2021).

2.1 | Notation

We denote natural numbers as $\mathbb{N}$ and n-dimensional Euclidean space as $\mathbb{R}^n$. Scalars are represented in lowercase italics (e.g., x), vectors are represented in lower case boldface (e.g., $\mathbf{x}$), and arrays and matrices are represented in uppercase boldface (e.g., $\mathbf{X}$).

For multiple vectors $\mathbf{x}_1, \cdots, \mathbf{x}_n$, we use the tuple notation to indicate the vertical concatenation of these vectors: $(\mathbf{x}_1, \cdots, \mathbf{x}_n) = [\mathbf{x}_1^\top, \cdots, \mathbf{x}_n^\top]^\top$. For a vector $\mathbf{v} \in \mathbb{R}^n$, we index its i-th element as $\mathbf{v}[i]$. For an array $\mathbf{A} \in \mathbb{R}^{n \times m}$, we index its i-th row (column) as $\mathbf{A}[i,:]$ ($\mathbf{A}[:,i]$). For a set $A = \{a(i)\}_{i=1}^{n_A} \subset \mathbb{R}^n$ ($n_A \in \mathbb{N} \cup \{\infty\}$), the notation $a(i)$ denotes the i-th element of the set.

Let $\mathbf{v}$ be the true value of the quantity of interest (e.g., a system state); we use a hat, $\hat{\mathbf{v}}$, to indicate an estimated value.

We use uppercase script characters to denote sets and set-valued functions. We use $\oplus$ to denote the Minkowski sum (i.e., set sum) operation; for a pair of sets $\mathcal{A}$ and $\mathcal{B}$, this operation is defined as $\mathcal{A} \oplus \mathcal{B} := \{a + b \mid a \in \mathcal{A}, b \in \mathcal{B}\}$. The negative of a set is denoted by $-\mathcal{A} = \{-a \mid a \in \mathcal{A}\}$.

2.2 | Probabilistic Zonotopes

We make use of zonotope-based SR analysis (Althoff, 2010; Althoff et al., 2009), which has seen recent success in integrity monitoring (Bhamidipati & Gao, 2020a, 2020b). A *zonotope* is a particular type of convex symmetrical polytope defined as follows:

$$\mathcal{Z}(\mathbf{c}, \mathbf{G}) = \{\mathbf{c} + \mathbf{G}\beta \in \mathbb{R}^n \mid \|\beta\|_\infty \leq 1\} \tag{1}$$

where $\mathbf{c} \in \mathbb{R}^n$ is the *center*, $\mathbf{G} \in \mathbb{R}^{n \times m}$ is a *generator matrix*, and $\beta \in \mathbb{R}^m$ is a *coefficient vector*. The columns of $\mathbf{G}$ are called *generators*. One can interpret a zonotope as the Minkowski sum of the center with a line segment created by scaling each generator by its corresponding coefficient, which lies in $[-1, +1]$. If $\mathcal{X} = \mathcal{Z}(\mathbf{c}, \mathbf{G})$, we denote $-\mathcal{X} = \mathcal{Z}(-\mathbf{c}, -\mathbf{G})$.

Zonotopes are bounded sets, which limits the types of distributions they can represent. To address this concern, Althoff et al. (2009) introduced *p-zonotopes* to represent an *enclosing probabilistic hull* (EPH), which is a conservative approximation of the set of estimated states and their corresponding probability distributions. These objects can be used to probabilistically overbound a collection of probability density functions, which one can obtain by performing reachability analysis on a system described by stochastic linear differential inclusions, as is considered in the present work.

Much like regular zonotopes, p-zonotopes are parameterized by a *center* and *generators*, which determine the zonotope *width*. For p-zonotopes, the width represents an uncertain mean of the underlying distributions. However, unlike a regular zonotope, a p-zonotope is additionally parameterized by a Gaussian *covariance* matrix. Thus, p-zonotopes represent an EPH of a set of distributions, e.g., a set of Gaussian distributions with a mean in the zonotope and the given covariance:

$$\mathcal{Z}_{\mathrm{p}}(\mathbf{c}, \mathbf{G}, \Sigma) = \{\mathbf{c} + \mathbf{G}\beta + \mathbf{w} \mid \|\beta\|_\infty \leq 1 \text{ and } \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \Sigma)\} \tag{2}$$

where $\Sigma \in \mathbb{R}^{n \times n}$ is a positive semi-definite covariance matrix. By modeling a probabilistic overbound on the state, p-zonotopes can encompass multiple possible Gaussian distributions, as shown in Figure 1, thereby allowing for uncertainty models with fewer restrictive assumptions on the error distribution (e.g., a zero-mean Gaussian distribution with perfectly known covariance).

Figure 1 depicts examples of 1D and 2D p-zonotopes. Note that a p-zonotope is not a probability distribution, but a probabilistic *overbound* over a set of probability distributions, also called an EPH. Also note that the 2D p-zonotope in Figure 1 provides a tighter probabilistic bound than one would obtain by maintaining two separate 1D overbounds, because of the relationship between the two states over

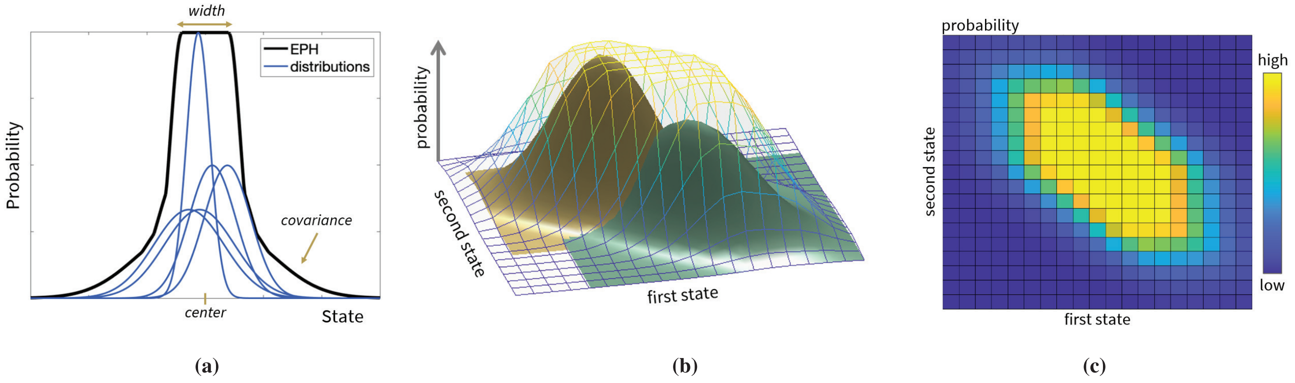

FIGURE 1 Illustrations of p-zonotopes
Note that a p-zonotope is not a probability distribution, but a probabilistic overbound over a set of state distributions, also called an EPH. **(a)** shows a 1D p-zonotope, represented by its *center*, *width*, and *covariance* parameters. **(b)** and **(c)** show a 2D example of a p-zonotope from a 3D view and a bird's eye view, respectively.

the distribution set. Although we depict p-zonotopes in 1D and 2D in Figure 1, these overbounding objects can be extended to any multidimensional space.

Note that p-zonotopes differ from mixture distributions in two key ways. First, these mathematical objects represent a probabilistic *overbound* on the set of possible distributions, rather than modeling a multimodal distribution with explicitly defined mixture weights. Second, mixture distributions often incorporate a finite number of distribution components or, potentially, a countably infinite number, whereas p-zonotopes can encompass an uncountably infinite number of possible distributions, which is especially useful for modeling certain types of distribution uncertainties. As an example, let us consider a Gaussian random variable that has an unknown, but bounded, mean. The set of possible distributions is uncountably infinite, but its overbound can be elegantly represented via a p-zonotope, as depicted in Figure 1.

Importantly, p-zonotopes are closed under linear maps and Minkowski sums (Althoff et al., 2009), i.e., the resulting mathematical object under these operations is also a p-zonotope. These two operations are necessary for reachability analysis, wherein a system's uncertain state is propagated forward under the system's dynamics. The resulting uncertain set of reachable states is typically dilated via the Minkowski sum to compensate for uncertainty (e.g., linearization error). Hence, p-zonotopes have recently been used for the verification of stochastic systems (e.g., as in Bhamidipati and Gao (2020a, 2020b) and Combastel and Zolghadri (2020)).

3 | PROPOSED METHOD

In this section, we first provide a high-level overview of our proposed method in Section 3.1. Next, we outline the derivation of the p-zonotope model for a general, loosely coupled Kalman filter that integrates GPS positioning measurements with odometry information from any self-contained sensor, such as an IMU. The odometry information from the self-contained sensor is incorporated through the filter propagation model, while the GPS positioning measurements are incorporated through the measurement model. Following a derivation similar to that of Shetty and Gao (2019), we first describe the point-valued Kalman filter-based state estimate in Section 3.2. Then, in Section 3.3, using an overbounding p-zonotope noise model of the process and measurement noises under authentic conditions, as well as the properties of linearity and Minkowski summation for p-zonotopes, we express the probabilistic overbound on the state estimation error for the SR-based Kalman filter (SR-KF). This derivation initially assumes a linear propagation model as represented by the self-contained sensor. In Section 3.4, we derive the corresponding SR-based Chimera spoofing statistic and detector. In Sections 3.5 and 3.6, we extend the derivations to a nonlinear propagation model.

Finally, we extend our derivation of the SR-based filter and Chimera spoofing detector for application to a nonlinear propagation model represented by a self-contained sensor, which applies to IMU body-frame measurements. We thereby derive an SR-based extended Kalman filter (SR-EKF) in Section 3.5 and the corresponding Chimera SR-EKF spoofing detector in Section 3.6.

3.1 | Overview

The core idea of our proposed method is as follows, with an illustration of our method shown in Figure 2 and the high-level architecture depicted in Figure 3. Recall that during the Chimera authentication period of $t_{\text{auth}} = 6$ s, the receiver

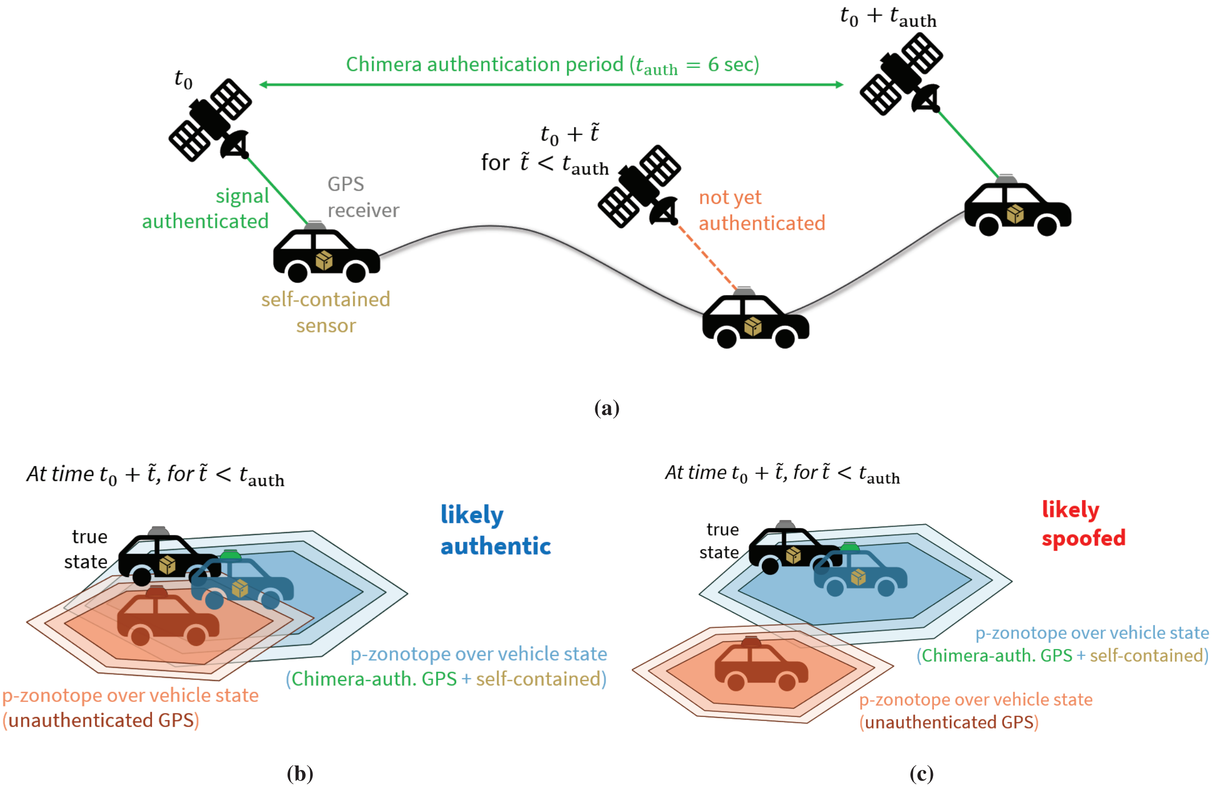
$t_0 + t_{\text{auth}}$
Chimera authentication period ($t_{\text{auth}} = 6$ sec)
t_0
$t_0 + \tilde{t}$
for $\tilde{t} < t_{\text{auth}}$
signal authenticated
GPS receiver
not yet authenticated
self-contained sensor
(a)
At time $t_0 + \tilde{t}$, for $\tilde{t} < t_{\text{auth}}$
likely authentic
true state
p-zonotope over vehicle state
(Chimera-auth. GPS + self-contained)
p-zonotope over vehicle state
(unauthenticated GPS)
(b)
At time $t_0 + \tilde{t}$, for $\tilde{t} < t_{\text{auth}}$
likely spoofed
true state
p-zonotope over vehicle state
(Chimera-auth. GPS + self-contained)
p-zonotope over vehicle state
(unauthenticated GPS)
(c)

FIGURE 2 Core idea and illustration of our proposed method (a) Illustration of the problem statement applied to a ground vehicle model, over one Chimera fast channel epoch of $t_{\text{auth}} = 6$ s (b) **Likely authentic** scenario, at time step $t_0 + \tilde{t}$ (c) **Likely spoofed** scenario, at time step $t_0 + \tilde{t}$

The time of the last Chimera authentication is represented as the first time step t_0, and $\tilde{t} < t_{\text{auth}}$ represents the duration of time between Chimera authentications, where the received GPS measurement is not yet authenticated, as depicted in **(a)**. In **(b)** and **(c)**, the p-zonotopes representing position state estimates depending on the received self-contained sensor information are shown in blue, and the corresponding p-zonotopes based on unauthenticated GPS measurements are shown in orange. The illustration in **(b)** depicts a scenario in which the received, unauthenticated GPS signal is likely authentic, while the illustration in **(c)** analogously depicts a scenario in which the unauthenticated GPS signal is likely spoofed, due to the inconsistency of the unauthenticated GPS-estimated p-zonotope over the vehicle state with respect to the p-zonotope estimated by the self-contained sensor.

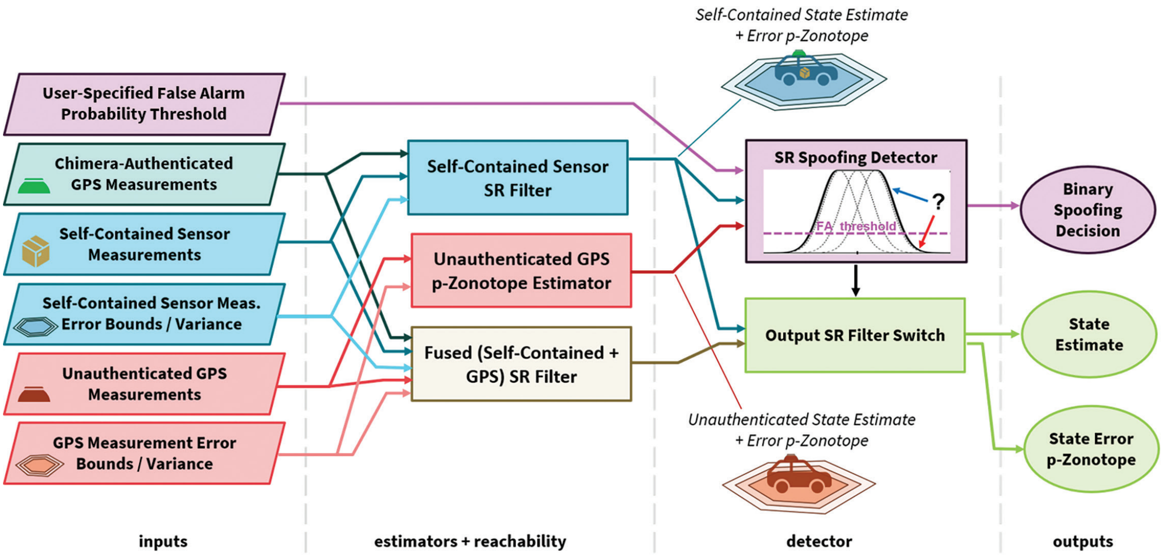
Self-Contained State Estimate
+ Error p-Zonotope
User-Specified False Alarm Probability Threshold
Chimera-Authenticated GPS Measurements
Self-Contained Sensor Measurements
Self-Contained Sensor Meas. Error Bounds / Variance
Unauthenticated GPS Measurements
GPS Measurement Error Bounds / Variance
Self-Contained Sensor SR Filter
Unauthenticated GPS p-Zonotope Estimator
Fused (Self-Contained + GPS) SR Filter
SR Spoofing Detector
?
FA threshold
Output SR Filter Switch
Binary Spoofing Decision
State Estimate
State Error p-Zonotope
Unauthenticated State Estimate
+ Error p-Zonotope
inputs
estimators + reachability
detector
outputs

FIGURE 3 High-level architecture of our proposed spoofing detector and SR estimator for continuous Chimera-enhanced GPS signal verification
Inputs are parallelograms, processes are boxes, and outputs are ellipses. Different subsets of the inputs are used in three separate SR state estimators: one with the self-contained sensor, one with unauthenticated GPS measurements, and one that fuses the self-contained sensor measurements with the unauthenticated GPS measurements. The SR state estimators output both state estimates and stochastic reachable sets, which are used to make a spoofing decision by the SR spoofing detector process box in violet, depending on whether the error between the self-contained sensor and unauthenticated GPS state estimates lies within a set of predicted errors. The final output of the detector is the binary spoofing decision, as well as the final state estimate and estimator error p-zonotope provided by the output SR filter switch in light green. If the spoofing detector outputs an "authentic" decision, the filter switch outputs the state estimate and p-zonotope from the fused SR filter; otherwise, the filter outputs the state estimate and p-zonotope from the self-contained SR filter. The 1D EPH image used for the SR spoofing detector process block is adapted from Althoff et al. (2009).

obtains a series of unauthenticated GPS positioning measurements, as shown in Figure 2(a). Between the Chimera authentication times, we maintain a pair of receiver position state estimates: one estimate is based on the unauthenticated GPS positioning measurements, and the other estimate is initialized according to the previous Chimera-authenticated GPS measurements and then updated according to a trusted local self-contained sensor, such as an IMU. We similarly maintain a pair of p-zonotopes on the receiver state error: one p-zonotope is based on the variance and bounded biases from the unauthenticated GPS positioning measurements during authentic conditions, and the other p-zonotope is computed via an SR-based state estimation filter based on the self-contained sensor. When initializing the SR filter at the beginning of the Chimera epoch, we assume that the user has access to a sufficient number of Chimera-enhanced GPS measurements from the previous epoch in order to obtain an authenticated position solution, e.g., at least four Chimera signals for a weighted least-squares solution.

From the two stochastic reachable sets, we find the probabilistic set, or EPH, of expected errors between the estimators, under nominal unspoofed conditions. To detect spoofing, we assess whether the current error between the estimators has a sufficiently high likelihood within this EPH with respect to a user-defined false alarm condition, as shown in the final detector block in Figure 3. Intuitively, if the received GPS signal is likely authentic, then we should observe significant overlap between the two p-zonotopes on the state estimate, as depicted in Figure 2(b). However, if the p-zonotope based on the unauthenticated GPS measurements is not sufficiently consistent with the p-zonotope based on the self-contained sensor information, as depicted in Figure 2(c), then we declare the received GPS measurements as being likely spoofed.

As shown in Figure 3, the output stochastic reachable state estimation takes the detector decision as an input. While the detector outputs an “authentic” decision, the output SR state estimator outputs the fused state estimate, based on the self-contained sensor measurement and the GPS positioning measurements. Once the detector outputs a “spoofed” decision, the output SR state estimator switches to rely on the self-contained sensor filter until it can re-authenticate the received GPS measurements via the Chimera enhancement.

3.2 | Point-Wise Kalman Filter Estimation Expressions

To establish the context for the SR-KF, we review the standard Kalman filter. Let us consider a generic linear receiver motion model:

$$\mathbf{x}_k = \mathbf{A}_k \mathbf{x}_{k-1} + \mathbf{B}_k \mathbf{u}_{k-1} + \Gamma_k \mathbf{w}_{k-1} \tag{3}$$

where $\mathbf{x}_k$, $\mathbf{u}_k$, and $\mathbf{w}_k$ represent the true state, state transition input, and process noise, respectively, at time k. In this work, we model the state propagation of the Kalman filter-based estimator by using the self-contained sensor information. For our point-wise state estimate, $\mathbf{w}_k$ is modeled according to the hypothesis that $\mathbf{w}_k \sim \mathcal{N}(0, \mathbf{W}_k)$. We discuss the extension to an unknown biased process noise, as modeled by a p-zonotope in Section 2.2, which allows us to express the probabilistic overbound for the state estimation errors. For our state estimate, we have the following Kalman filter expressions for the prediction step:

$$\hat{\mathbf{x}}_{k|k-1} = \mathbf{A}_k \hat{\mathbf{x}}_{k-1} + \mathbf{B}_k \mathbf{u}_{k-1} \tag{4}$$

$$\hat{\mathbf{P}}_{k|k-1} = \mathbf{A}_k \hat{\mathbf{P}}_{k-1} \mathbf{A}_k^\top + \boldsymbol{\Gamma}_k \mathbf{W}_{k-1} \boldsymbol{\Gamma}_k^\top \tag{5}$$

where $\hat{\mathbf{x}}_k$ and $\hat{\mathbf{x}}_{k|k-1}$ represent the state estimate and predicted state estimate, respectively, at time k. Analogously, $\hat{\mathbf{P}}_k$ and $\hat{\mathbf{P}}_{k|k-1}$ represent the covariance of the state estimate and predicted state estimate, respectively, at time k. For the Kalman filter update step, which corrects the predicted state using the latest received GPS positioning measurements $\mathbf{z}_k$, we have the following:

$$\mathbf{z}_k = \mathbf{H}_k \mathbf{x}_k + \mathbf{r}_k \tag{6}$$

$$\mathbf{K}_k = \hat{\mathbf{P}}_{k|k-1} \mathbf{H}_k^\top \left(\mathbf{H}_k \hat{\mathbf{P}}_{k|k-1} \mathbf{H}_k^\top + \mathbf{Z}_k \right)^{-1} \tag{7}$$

$$\hat{\mathbf{x}}_k = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k \left(\mathbf{z}_k - \mathbf{H}_k \hat{\mathbf{x}}_{k|k-1} \right) \tag{8}$$

$$\hat{\mathbf{P}}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \hat{\mathbf{P}}_{k|k-1} \tag{9}$$

where $\mathbf{H}_k$ is the system measurement matrix for the GPS positioning measurements, $\mathbf{r}_k$ is the measurement noise, and $\mathbf{K}_k$ is the Kalman gain matrix. For our point-wise state estimate, the measurement noise is modeled according to the hypothesis that $\mathbf{r}_k \sim \mathcal{N}(0, \mathbf{Z}_k)$, where $\mathbf{Z}_k$ is the measurement covariance. We discuss the extension to an unknown biased measurement noise vector in Section 3.3.

3.3 | Stochastic Reachability-Based Kalman Filter (SR-KF)

Now, by following a procedure similar to that of Shetty and Gao (2019), we derive the stochastic reachable set of state estimation errors in the presence of unknown bounded biases in the process noise and measurement noise. First, we represent the state estimation error at time k as follows:

$$\tilde{\mathbf{x}}_k = \hat{\mathbf{x}}_k - \mathbf{x}_k \tag{10}$$

Substituting $\mathbf{x}_k$ with the expression in Equation (3) and $\hat{\mathbf{x}}_k$ with the expressions in Equations (8) and (4), we can write:

$$\tilde{\mathbf{x}}_k = \left(\mathbf{A}_k\hat{\mathbf{x}}_{k-1} + \mathbf{B}_k\mathbf{u}_{k-1} + \mathbf{K}_k\left(\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_{k|k-1}\right)\right) - \left(\mathbf{A}_k\mathbf{x}_{k-1} + \mathbf{B}_k\mathbf{u}_{k-1} + \boldsymbol{\Gamma}_k\mathbf{w}_{k-1}\right) \tag{11}$$

$$= \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} + \mathbf{K}_k\left(\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_{k|k-1}\right) - \boldsymbol{\Gamma}_k\mathbf{w}_{k-1} \tag{12}$$

Then, using Equations (6), (4), and (3), we derive a recursive expression for the state estimation error:

$$\tilde{\mathbf{x}}_k = \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} + \mathbf{K}_k\left(\mathbf{H}_k\mathbf{x}_k + \mathbf{r}_k - \mathbf{H}_k\left(\mathbf{A}_k\hat{\mathbf{x}}_{k-1} + \mathbf{B}_k\mathbf{u}_{k-1}\right)\right) - \boldsymbol{\Gamma}_k\mathbf{w}_{k-1} \tag{13}$$

$$\begin{aligned} &= \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} + \mathbf{K}_k\left(\mathbf{H}_k\left(\mathbf{A}_k\mathbf{x}_{k-1} + \mathbf{B}_k\mathbf{u}_{k-1} + \boldsymbol{\Gamma}_k\mathbf{w}_{k-1}\right) + \mathbf{r}_k\right.\\ &\quad \left.-\mathbf{H}_k\left(\mathbf{A}_k\hat{\mathbf{x}}_{k-1} + \mathbf{B}_k\mathbf{u}_{k-1}\right)\right) - \boldsymbol{\Gamma}_k\mathbf{w}_{k-1} \end{aligned} \tag{14}$$

$$= \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} + \mathbf{K}_k\left(-\mathbf{H}_k\mathbf{A}_k\tilde{\mathbf{x}}_{k-1} + \mathbf{H}_k\Gamma_k\mathbf{w}_{k-1} + \mathbf{r}_k\right) - \Gamma_k\mathbf{w}_{k-1} \tag{15}$$

$$= \left(\mathbf{I} - \mathbf{K}_k\mathbf{H}_k\right)\mathbf{A}_k\tilde{\mathbf{x}}_{k-1} - \left(\mathbf{I} - \mathbf{K}_k\mathbf{H}_k\right)\Gamma_k\mathbf{w}_{k-1} + \mathbf{K}_k\mathbf{r}_k \tag{16}$$

Converting Equation (16) to set notation as is done in Shetty and Gao (2019), we obtain a recursive expression for the stochastic set of state errors $\tilde{\mathcal{X}}$, which represents a probabilistic overbound on the set of state errors:

$$\tilde{\mathcal{X}}_k = \left(\left(\mathbf{I} - \mathbf{K}_k\mathbf{H}_k\right)\mathbf{A}_k\tilde{\mathcal{X}}_{k-1}\right) \oplus \left(\left(-\mathbf{I} + \mathbf{K}_k\mathbf{H}_k\right)\Gamma_k\mathcal{W}_{k-1}\right) \oplus \left(\mathbf{K}_k\mathcal{R}_k\right) \tag{17}$$

where $\mathcal{W}_k$ and $\mathcal{R}_k$ represent the stochastic reachable sets of errors in the process noise and measurement noise, respectively. Finally, we can represent a probabilistic overbound over the set of true states, in terms of the state estimate and stochastic set of state errors:

$$\mathcal{X}_k = \hat{\mathbf{x}}_k + \left(-\tilde{\mathcal{X}}_k\right) \tag{18}$$

3.4 | Chimera SR-KF Spoofing Detector

Next, we leverage the SR-KF for spoofing detection. In particular, during the time interval between Chimera authentications, e.g., 1.5-s or 6-s intervals proposed for the Chimera fast channel implementations, we separately model the receiver state estimates from the self-contained sensor information using the state propagation model described in Equation (4) as follows:

$$\hat{\mathbf{x}}_k^{\text{self}} = \mathbf{A}_k\hat{\mathbf{x}}_{k-1}^{\text{self}} + \mathbf{B}_k\mathbf{u}_{k-1} \tag{19}$$

where the superscript $(\cdot)^{\text{self}}$ indicates self-contained sensor information. The corresponding state error is similarly derived as in Equation (12), but without the measurement update term, thereby allowing us to model the stochastic set of state errors as follows:

$$\tilde{\mathcal{X}}_k^{\text{self}} = \mathbf{A}_k \tilde{\mathcal{X}}_{k-1}^{\text{self}} \oplus \left(-\Gamma_k\right)\mathcal{W}_{k-1} \tag{20}$$

Projecting the state $\hat{\mathbf{x}}_k^{\text{self}}$ to the position coordinates to obtain $\hat{\mathbf{p}}_k^{\text{self}}$, we define the spoofing statistic $\mathbf{q}_k$ as the difference in position state estimates:

$$\mathbf{q}_k := \hat{\mathbf{p}}_k^{\text{self}} - \hat{\mathbf{p}}_k^{\text{GPS}} \tag{21}$$

where $\hat{\mathbf{p}}_k^{\text{GPS}}$ represents the received GPS positioning measurement. Next, we can relate the spoofing statistic to the difference in position state errors in the following way:

$$\mathbf{q}_k = \hat{\mathbf{p}}_k^{\text{self}} - \hat{\mathbf{p}}_k^{\text{GPS}} \tag{22}$$

$$= \left(\mathbf{p}_k + \tilde{\mathbf{p}}_k^{\text{self}}\right) - \left(\mathbf{p}_k + \tilde{\mathbf{p}}_k^{\text{GPS}}\right) \tag{23}$$

$$= \tilde{\mathbf{p}}_k^{\text{self}} - \tilde{\mathbf{p}}_k^{\text{GPS}} \tag{24}$$

As a result, assuming that the state errors from the self-contained sensor noise and the GPS noise are independent, we can model the stochastic set of the spoofing statistic during nominal conditions as follows:

$$\mathcal{Q}_k = \tilde{\mathcal{P}}_k^{\text{self}} \oplus \left(-\tilde{\mathcal{P}}_k^{\text{GPS}}\right) \tag{25}$$

where $\tilde{\mathcal{P}}_k^{\text{self}}$ is the stochastic set of position errors found by projecting $\tilde{\mathcal{X}}_k^{\text{self}}$ to the position coordinates and $\tilde{\mathcal{P}}_k^{\text{GPS}}$ is the stochastic set of errors for the GPS positioning measurement $\hat{\mathbf{p}}_k^{\text{GPS}}$, modeled according to the variance and bounded biases from the GPS positioning measurements under nominal unspoofed conditions. Because these stochastic sets represent a probabilistic overbound on the set of state errors under nominal conditions, the stochastic set of $\mathcal{Q}_k$ correspondingly represents a probabilistic overbound on the spoofing statistic in the nominal case.

Let $\mathcal{Q}_k(\mathbf{q}_k)$ represent the evaluation of the p-zonotope $\mathcal{Q}_k$ at the spoofing statistic $\mathbf{q}_k$, which corresponds to the probabilistic overbound of the spoofing statistic, under nominal conditions. Thus, we choose the binary spoofing decision to satisfy a user-defined false alarm requirement P_{FA} through the following definition:

$$d_k := \mathbb{1}\left\{\mathcal{Q}_k(\mathbf{q}_k) \geq P_{\text{FA}}\right\} \tag{26}$$

where $\mathbb{1}$ returns 1 if its argument is true and 0 if false. Thus, $d_k = 1$ corresponds to an "authentic" decision, and $d_k = 0$ corresponds to a "spoofed" decision, which are chosen according to the user-defined false alarm probability. Intuitively, if the probability of the spoofing statistic $\mathbf{q}_k$ under the nominal condition assumptions represented by the p-zonotope $\mathcal{Q}_k$ is sufficiently large, we declare that the signal is likely authentic. Otherwise, we declare that the signal is likely spoofed. The threshold of P_{FA} guarantees that we satisfy the corresponding user-defined false alarm probability requirement under authentic conditions.

3.5 | Extension to a Nonlinear Propagation Model: Stochastic Reachability-Based Extended Kalman Filter (SR-EKF)

To estimate the state error p-zonotopes and perform validation with a self-contained sensor with a nonlinear model, such as an IMU that returns accelerations in the body frame of reference, we must extend the SR-KF and spoofing detector for application to a nonlinear propagation model **f**, which we represent as follows:

$$\mathbf{x}_k = \mathbf{f}\left(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}, \mathbf{w}_{k-1}\right) \tag{27}$$

Maintaining our measurement model of GPS positioning measurements from Section 3.3, we only need to update the predicted state estimate expressions from Equation (4) to obtain our point-wise EKF estimate:

$$\hat{\mathbf{x}}_{k|k-1} = \mathbf{f}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right) \tag{28}$$

Here, the remaining point-wise state estimation expressions from Section 3.3 remain the same for our application for the nonlinear propagation model extension. We define $\mathbf{A}_k$ and Γ_k as follows for the SR-EKF:

$$\mathbf{A}_k := \frac{\partial \mathbf{f}}{\partial \mathbf{x}}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right) \quad \text{and} \quad \Gamma_k := \frac{\partial \mathbf{f}}{\partial \mathbf{w}}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right) \tag{29}$$

To derive the stochastic set of state estimation errors for the SR-EKF, we must account for linearization error in the state estimation expression. To do so, we follow the method in Althoff et al. (2008) to conservatively represent the Lagrange remainder of the nonlinear propagation expression. To proceed, we define $\mathbf{s}_k$ as the concatenation of all arguments of the nonlinear propagation model **f** at time step *k*:

$$\mathbf{s}_k := \left(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k\right) \tag{30}$$

Correspondingly, we can express the true state propagation as $\mathbf{x}_k = \mathbf{f}\left(\mathbf{s}_{k-1}\right)$. We can further write the true state propagation in terms of the first-order approximation of the Taylor series along with the corresponding Lagrange remainder term $\ell^{\hat{\mathbf{s}}}_{k-1}$, as follows:

$$\mathbf{x}_k = \mathbf{f}\left(\hat{\mathbf{s}}_{k-1}\right) + \left.\frac{\partial \mathbf{f}(\mathbf{s})}{\partial \mathbf{s}}\right|_{\mathbf{s}=\hat{\mathbf{s}}_{k-1}} \left(\mathbf{s}_{k-1} - \hat{\mathbf{s}}_{k-1}\right) + \ell^{\hat{\mathbf{s}}}_{k-1} \tag{31}$$

$$\ell^{\hat{\mathbf{s}}}_{k-1} := \frac{1}{2}\left(\mathbf{s}_{k-1} - \hat{\mathbf{s}}_{k-1}\right)^{\top} \frac{\partial^2 \mathbf{f}\left(\zeta\right)}{\partial \mathbf{s}^2}\left(\mathbf{s}_{k-1} - \hat{\mathbf{s}}_{k-1}\right) \tag{32}$$

where $\hat{\mathbf{s}}_k := \begin{bmatrix} \hat{\mathbf{x}}_k^{\top} & \mathbf{u}_k^{\top} & 0 \end{bmatrix}^{\top}$. Note that $\mathbf{s}_k$ is restricted to be in a convex set, because we represent each component as a p-zonotope. As a result, for a particular $\mathbf{s}_k$ and $\hat{\mathbf{s}}_k$, we have $\zeta \in \{\hat{\mathbf{s}}_k + \alpha\left(\mathbf{s}_k - \hat{\mathbf{s}}_k\right) \mid \alpha \in [0,1]\}$ (Althoff et al., 2008; Berz & Hoffstätter, 1998).

Rewriting Equation (31) in terms of each component of $\mathbf{s}$, we can express the true state propagation in terms of the linearized SR-EKF matrices $\mathbf{A}_k$, Γ_k, and $\mathbf{B}_k := \frac{\partial \mathbf{f}}{\partial \mathbf{u}}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right)$, with:

$$\mathbf{x}_k = \mathbf{f}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right) + \mathbf{A}_k\left(\mathbf{x}_{k-1} - \hat{\mathbf{x}}_{k-1}\right) + \mathbf{B}_k\left(\mathbf{u}_{k-1} - \mathbf{u}_{k-1}\right) + \Gamma_k\left(\mathbf{w}_{k-1} - 0\right) + \ell^{\hat{\mathbf{s}}}_{k-1} \tag{33}$$

$$= \mathbf{f}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right) - \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} + \Gamma_k\mathbf{w}_{k-1} + \ell^{\hat{\mathbf{s}}}_{k-1} \tag{34}$$

Correspondingly, for the SR-EKF, we can express the state error from Equation (10) by replacing the nonlinear predicted state representation with Equation (28) and by substituting the expression from Equation (34):

$$\tilde{\mathbf{x}}_k = \hat{\mathbf{x}}_k - \mathbf{x}_k \tag{35}$$

$$= \left(\hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k\left(\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_{k|k-1}\right)\right) - \left(\mathbf{f}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right) - \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} + \mathbf{\Gamma}_k\mathbf{w}_{k-1} + \ell^{\hat{\mathbf{s}}}_{k-1}\right) \tag{36}$$

$$\begin{aligned} &= \left(\mathbf{f}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right) + \mathbf{K}_k\left(\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_{k|k-1}\right)\right) - \left(\mathbf{f}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right)\right. \\ &\quad \left. - \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} + \mathbf{\Gamma}_k\mathbf{w}_{k-1} + \ell^{\hat{\mathbf{s}}}_{k-1}\right) \end{aligned} \tag{37}$$

$$= \mathbf{K}_k\left(\mathbf{z}_k - \mathbf{H}_k\mathbf{f}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right)\right) + \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} - \mathbf{\Gamma}_k\mathbf{w}_{k-1} - \ell^{\hat{\mathbf{s}}}_{k-1} \tag{38}$$

Using Equation (6) and replacing $\mathbf{x}_k$ with Equation (34), we finally derive a recursive expression of the state estimation error:

$$\tilde{\mathbf{x}}_k = \mathbf{K}_k\left(\mathbf{H}_k\mathbf{x}_k + \mathbf{r}_k - \mathbf{H}_k\mathbf{f}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right)\right) + \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} - \mathbf{\Gamma}_k\mathbf{w}_{k-1} - \ell^{\hat{\mathbf{s}}}_{k-1} \tag{39}$$

$$\begin{aligned} &= \mathbf{K}_k\left(\mathbf{H}_k\left(\mathbf{f}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right) - \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} + \mathbf{\Gamma}_k\mathbf{w}_{k-1} + \ell^{\hat{\mathbf{s}}}_{k-1}\right) + \mathbf{r}_k - \mathbf{H}_k\mathbf{f}\left(\hat{\mathbf{x}}_{k-1}, \mathbf{u}_{k-1}, 0\right)\right) \\ &\quad + \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} - \mathbf{\Gamma}_k\mathbf{w}_{k-1} - \ell^{\hat{\mathbf{s}}}_{k-1} \end{aligned} \tag{40}$$

$$= \mathbf{K}_k\left(\mathbf{H}_k\left(-\mathbf{A}_k\tilde{\mathbf{x}}_{k-1} + \mathbf{\Gamma}_k\mathbf{w}_{k-1} + \ell^{\hat{\mathbf{s}}}_{k-1}\right) + \mathbf{r}_k\right) + \mathbf{A}_k\tilde{\mathbf{x}}_{k-1} - \mathbf{\Gamma}_k\mathbf{w}_{k-1} - \ell^{\hat{\mathbf{s}}}_{k-1} \tag{41}$$

$$= \left(\mathbf{I} - \mathbf{K}_k\mathbf{H}_k\right)\mathbf{A}_k\tilde{\mathbf{x}}_{k-1} - \left(\mathbf{I} - \mathbf{K}_k\mathbf{H}_k\right)\mathbf{\Gamma}_k\mathbf{w}_{k-1} - \left(\mathbf{I} - \mathbf{K}_k\mathbf{H}_k\right)\ell^{\hat{\mathbf{s}}}_{k-1} + \mathbf{K}_k\mathbf{r}_k \tag{42}$$

Converting Equation (42) to set notation, we obtain a recursive expression for the stochastic set of state errors for the SR-EKF that incorporates the nonlinear propagation model:

$$\tilde{\mathcal{X}}_k = \left(\mathbf{I} - \mathbf{K}_k\mathbf{H}_k\right)\mathbf{A}_k\tilde{\mathcal{X}}_{k-1} \oplus \left(-\mathbf{I} + \mathbf{K}_k\mathbf{H}_k\right)\mathbf{\Gamma}_k\mathcal{W}_{k-1} \oplus \left(-\mathbf{I} + \mathbf{K}_k\mathbf{H}_k\right)\mathcal{L}^{\hat{\mathbf{s}}}_{k-1} \oplus \mathbf{K}_k\mathcal{R}_k \tag{43}$$

where $\mathcal{L}^{\hat{\mathbf{s}}}_{k-1}$ is a set representing all possible Lagrange remainder values $\ell^{\hat{\mathbf{s}}}_{k-1}$ given by choosing $\mathbf{s}_k$, $\hat{\mathbf{s}}_k$, and ζ as noted above. We overapproximate $\mathcal{L}^{\hat{\mathbf{s}}}_{k-1}$ as a zonotope using the strategy in Althoff et al. (2008). First, note that from Equation (10) and the definition of $\hat{\mathbf{s}}_k$, the state estimation error is contained in a set as follows:

$$\left(\mathbf{s}_k - \hat{\mathbf{s}}_k\right) \in \left(-\tilde{\mathcal{X}}_k\right) \times \left\{\mathbf{0}_{m\times 1}\right\} \times \mathcal{W}_k \tag{44}$$

where the set on the right-hand side is a p-zonotope created by the Cartesian product of the state, input, and noise p-zonotopes and m is the dimension of the input vector. By choosing a confidence level, we can overapproximate this set with a zonotope. Then, the method in Althoff et al. (2008) overapproximates this zonotope as an interval and finally overapproximates $\mathcal{L}^{\hat{\mathbf{s}}}_{k-1}$ as a zonotope by evaluating Equation (32) with interval arithmetic.

3.6 | Extension to a Nonlinear Propagation Model: Chimera SR-EKF Spoofing Detector

When using a nonlinear propagation model, we model the state estimate based on the self-contained sensor as follows:

$$\hat{\mathbf{x}}^{\text{self}}_k = \mathbf{f}\left(\hat{\mathbf{x}}^{\text{self}}_{k-1}, \mathbf{u}_{k-1}, 0\right) \tag{45}$$

where we similarly obtain the position estimate $\hat{\mathbf{p}}_k^{\text{self}}$ by extracting the corresponding position states. The corresponding state error can be similarly derived as in Equation (38) without the measurement update term, thereby allowing us to model the stochastic set of state errors as follows:

$$\tilde{\mathcal{X}}_k^{\text{self}} = \mathbf{A}_k^{\text{self}} \tilde{\mathcal{X}}_{k-1}^{\text{self}} \oplus \left(-\boldsymbol{\Gamma}_k^{\text{self}}\right) \mathcal{W}_{k-1} \oplus \left(-\mathcal{L}_{k-1}^{\hat{\mathbf{s}}^{\text{self}}}\right) \tag{46}$$

where $\hat{\mathbf{s}}_k^{\text{self}} := \left(\hat{\mathbf{x}}_k^{\text{self}}, \mathbf{u}_k, 0\right)$ and where $\mathbf{A}_k^{\text{self}}$ and $\boldsymbol{\Gamma}_k^{\text{self}}$ are defined similar to the definitions in Equation (29), but with respect to $\hat{\mathbf{s}}_k^{\text{self}}$. Similar to the linear SR-KF case, the corresponding p-zonotope on the position estimation error $\tilde{\mathcal{P}}_k^{\text{self}}$ can be found by projecting the error zonotope $\tilde{\mathcal{X}}_k^{\text{self}}$ onto the position domain. Using the updated p-zonotope for the nonlinear dynamics, the spoofing statistic for the SR-EKF $\mathbf{q}_k$ is the same as in Equation (24) whereas the corresponding stochastic set $\mathcal{Q}$ is the same as in Equation (25).

4 | EXPERIMENTAL RESULTS

In this section, we present two experimental validations of our proposed Chimera spoofing detector. We first implement and validate an example of the linear Chimera SR-KF detector and estimator in Section 3.4, followed by an example of a nonlinear extension of the Chimera SR-EKF detector and estimator in Section 3.6. We assume a 6-s fast channel implementation of Chimera and start our simulation at the beginning of the Chimera epoch, when the GPS signals can be authenticated. We assume that the user has access to a sufficient number of Chimera measurements to obtain an authenticated position solution at the start of the epoch, which our SR-KF and SR-EKF leverage for initialization, as indicated in Figure 3.

In both experimental examples, we run Monte Carlo simulations for 1000 sampled trajectories to probabilistically validate that our proposed detector maintains the required correct authentication rate (CAR), derived from the required false alarm rate (FAR), under authentic conditions. To evaluate the CAR, at each time step, we examine the ratio of binary spoofing decisions that performed a correct authentication across the 1000 Monte Carlo simulations. We similarly quantify the missed detection rate (MDR) and, consequently, the correct detection rate (CDR) in the presence of simulated trajectory-drifting spoofing attacks. Because our proposed SR-KF detector and estimator analyze discrepancies between the GPS and self-contained sensor measurements within the position domain, we similarly perform the simulated attacks in the position domain by modeling the spoofed GPS measurements as additive biases with respect to the true vehicle state. We additionally plot the ratio of states bounded by the 3σ confidence-level zonotopes from our Chimera SR estimator, during each scenario. Furthermore, to better understand when our Chimera SR detector and estimator switch between the fused SR filter and the self-contained SR filter, we additionally plot the 3σ state bounding statistics for the naively fused SR filter, for reference. To perform p-zonotope operations and assess the SR, we use the MATLAB CORA toolbox[1] (Althoff et al., 2018).

[1] We used the 2020 version of CORA, available at https://tumcps.github.io/CORA/.

4.1 | Validation of Chimera SR-KF Detector and Estimator with a Double-Integrator System

4.1.1 | *Setup*

To validate the Chimera SR-KF detector and estimator, we consider a linear double-integrator system model for a ground vehicle, with 2D inertial-frame accelerations as the self-contained sensor information, which is utilized for the vehicle propagation model. We assume that the propagation occurs with bounded biases within $[-0.1, +0.1]$ m in each position dimension and within $[-0.01, +0.01]$ ms^{-1} in each velocity dimension and with stochastic noise represented by the following covariance matrix:

$$Q = 0.1\ \mathrm{m}^3\mathrm{s}^{-3}\begin{bmatrix} \frac{t_s^3}{3}\mathbf{I}_2 & \frac{t_s^2}{2}\mathbf{I}_2 \\ \frac{t_s^2}{2}\mathbf{I}_2 & t_s\mathbf{I}_2 \end{bmatrix} \tag{47}$$

where t_s is the discrete sample period.

We additionally simulate unauthenticated GPS positioning measurements such that, in the nominal case, the measurements $\mathbf{z}^{\mathrm{GPS}}$ have a bounded bias of ±0.5 m in each position dimension with a standard deviation of 5 m. In the spoofed case, the measurements $\mathbf{z}^{\mathrm{GPS}}$ contain an additive ramping bias, which results in a total error of 60 m at the end of the simulated trajectory. The self-contained sensor and GPS measurements are both simulated at a rate of 10 Hz.

4.1.2 | *Results and Discussion*

For nominal unspoofed conditions, we observe in Figure 4(a) that the CAR of the Chimera SR-KF consistently remains above the 3σ confidence level of 0.997, thereby satisfying the corresponding user-specified FAR requirement of 0.003, where $\mathrm{CAR} = 1 - \mathrm{FAR}$. The continuously authentic decision is qualitatively validated by the bird's eye view of the trajectory in Figure 4(b), which shows the significant overlap of the stochastic sets $\tilde{\mathcal{P}}_k^{\mathrm{GPS}}$ and $\tilde{\mathcal{P}}_k^{\mathrm{self}}$. In Figure 4(b), the 3σ error zonotopes are plotted to be centered about the average estimated trajectory across the 1000 Monte Carlo runs for ease of visual interpretation, whereas the SR-KF centers the error zonotopes about the current state estimate in practice and for each of the conducted Monte Carlo simulations, as also indicated in Equation (18). Thus, the SR-KF correspondingly provides the user with a probabilistic bound on the true vehicle state. We further analyze this probabilistic bound on the true vehicle state in Figure 4(c), which plots the ratio of true trajectories bounded by our proposed Chimera SR-KF across the 1000 Monte Carlo simulations. The Chimera SR-KF uses state estimates and error p-zonotopes from the fused SR-KF when a spoofing event is not detected. As a result, in this nominal scenario, both the Chimera SR-KF and the naively fused SR-KF nearly always output the same state estimate and error p-zonotopes, which bound the true state for all Monte Carlo simulations in this case.

During the simulated spoofing attack, we observe in Figure 5(a) a low CDR in the initial part of the trajectory, when the bias is too small with respect to the expected nominal GPS measurement errors and self-contained sensor errors to be detected. Correspondingly, we observe a high MDR in the initial part of the trajectory. Once the bias is sufficiently large, the CDR increases to 1 and the MDR decreases to 0. Due to the presence of measurement biases, the difference in position estimate $\mathbf{q}_k$ lies

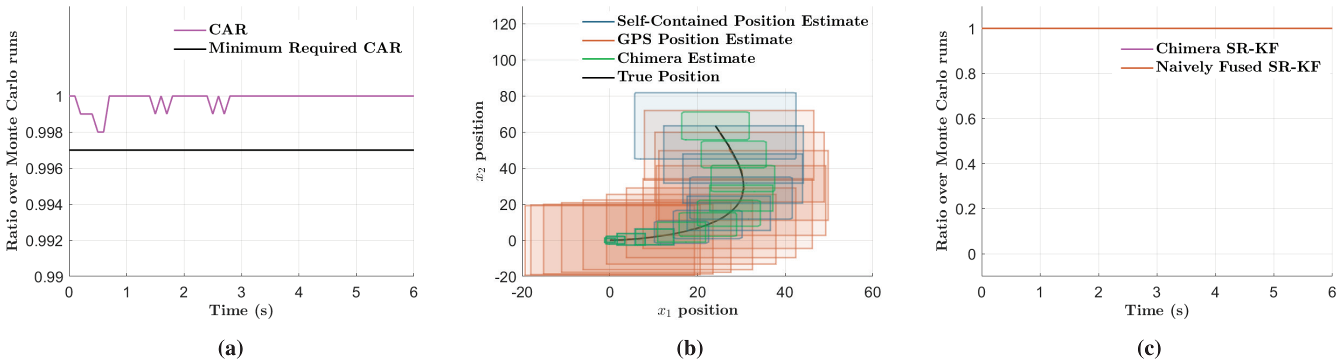

(a)

(b)

(c)

FIGURE 4 Monte Carlo validation of the Chimera SR-KF estimator and detector for the linear double-integrator system outlined in Section 4.1.1 under nominal conditions (a) Ratio of CAR (b) Bird's eye view of the trajectory (c) Ratio of states bound by 3σ error zonotopes
Trajectories and statistics have been averaged over 1000 Monte Carlo simulations for the same trajectory. In **(a)**, we observe that the CAR remains consistently above the 3σ confidence level, corresponding to a FAR of 0.003. We qualitatively observe in **(b)** that the 3σ confidence level zonotopes of $\tilde{\mathcal{P}}_k^{\text{GPS}}$ and $\tilde{\mathcal{P}}_k^{\text{self}}$ significantly overlap each other throughout the trajectory, indicating that the received GPS measurement is likely authentic. For ease of visual interpretation, the error zonotopes are plotted to be centered about the average estimated trajectory across the 1000 Monte Carlo runs, which also coincides with the true trajectory. In **(c)**, because the Chimera SR-KF state estimate and error p-zonotopes are nearly always identical to those of the naively fused SR-KF, we observe that both estimators consistently bound the true state over the trajectory.

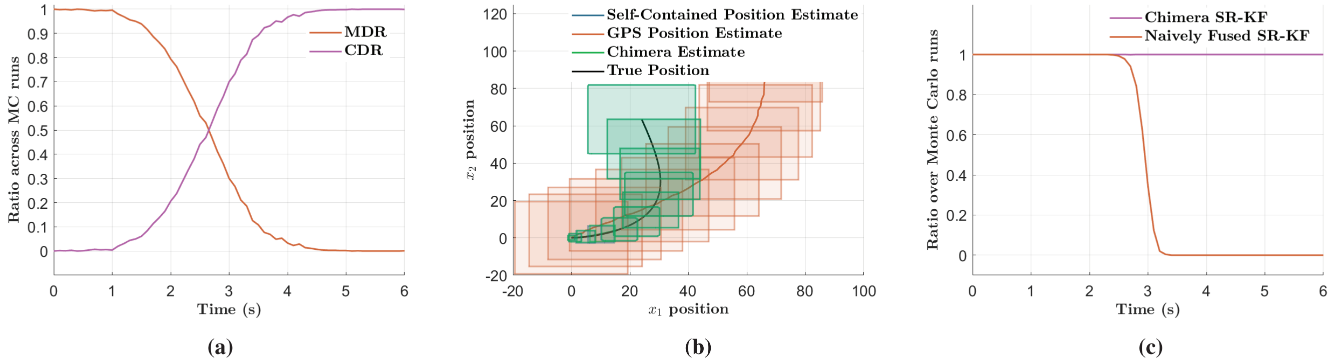

(a)

(b)

(c)

FIGURE 5 Monte Carlo (MC) validation of the Chimera SR-KF estimator and detector during spoofed conditions for the linear double-integrator system outlined in Section 4.1.1 (a) Ratio of MDR and CDR (b) Bird's eye view of the trajectory (c) Ratio of states bound by 3σ error zonotopes
Trajectories and statistics have been averaged over 1000 Monte Carlo simulations of the same trajectory. **(a)** We observe that the CDR increases to 1, while the MDR correspondingly decreases to 0, once the bias is sufficiently large to be detected by our proposed detector. **(b)** We observe an overlap between the self-contained and GPS 3σ error zonotopes in the initial part of the trajectory, resulting in missed detections of the spoofing event. For ease of visual interpretation, the error zonotopes are plotted to be centered about the average Chimera SR-KF and GPS estimated trajectories across the 1000 Monte Carlo runs, where the average Chimera SR-KF estimated trajectory also coincides with the true trajectory. As the spoofed trajectory deviates from the true trajectory, our proposed detector detects the spoofing event and the Chimera SR-KF switches to use the self-contained SR filter estimate. Thus, the Chimera SR-KF error zonotopes become equivalent to those of the self-contained SR-KF and correspondingly overlap exactly for the latter half of the trajectory. **(c)** We observe that in the initial part of the trajectory, both the Chimera SR-KF and naively fused SR-KF bound the true state. In the latter part of the trajectory, the Chimera SR-KF continues to bound the true state while the naively fused SR-KF stops bounding the true state on average.

outside the set $\mathcal{Q}_k$ for a sufficiently large bias, thereby causing the detector to raise an alarm when the bias is sufficiently large. Qualitatively, we correspondingly observe a limited area of intersection between the stochastic sets $\tilde{\mathcal{P}}_k^{\text{GPS}}$ and $\tilde{\mathcal{P}}_k^{\text{self}}$ in the bird's eye view of the trajectory in Figure 5(b).

Once a spoofing event is detected, the Chimera SR-KF switches from using the fused state estimates and error p-zonotopes to the self-contained state estimates and error p-zonotopes. When the GPS measurement bias is small, the fused state estimate and 3σ error zonotope bound the true state. As the bias grows, the spoofing attack is detected by our proposed approach, and eventually, the fused state estimate and 3σ zonotopes no longer bound the true state. In this case, the Chimera SR-KF switches to using the state estimates and error p-zonotopes of the self-contained SR filter, and we observe in Figure 5(c) that the Chimera SR-KF continues to bound the true state during this spoofing scenario over the 1000 Monte Carlo trajectories.

4.2 | Validation of Chimera SR-EKF Detector and Estimator with an IMU-GPS System

4.2.1 | Setup

To validate the Chimera SR-EKF detector and estimator with a nonlinear state propagation model, we model an IMU sensor as the self-contained sensor, with stochastic measurement noise and bounded measurement biases. We model the vehicle state as $\mathbf{x} = \left[\mathbf{p}^\top, \mathbf{v}^\top, \psi\right]^\top$ where $\mathbf{p}, \mathbf{v} \in \mathbb{R}^2$ are the 2D position and velocity states, respectively, and ψ is the heading angle. The vehicle's state propagation is modeled as a nonlinear system:

$$\mathbf{p}_k = \mathbf{p}_{k-1} + t_\text{s}\mathbf{v}_{k-1}, \tag{48}$$

$$\mathbf{v}_k = \mathbf{v}_{k-1} + t_\text{s}\mathbf{R}_{k-1}^{\text{ref}}\mathbf{a}_{k-1} \tag{49}$$

$$\psi_k = \psi_{k-1} + t_\text{s}\omega_{k-1} \tag{50}$$

where t_s is the sample period. In the above model, $\mathbf{R}_{k-1}^{\text{ref}}$ denotes the passive rotation matrix from the body frame of reference to the inertial frame of reference and is the source of nonlinearity in the propagation model. The 2D body-frame accelerations $\mathbf{a}_{k-1}$ and angular velocity ω_{k-1} are noisy measurements with unknown bounded biases from the IMU sensor and are utilized to propagate the system state.

We simulate both the IMU measurements and GPS positioning measurements at a rate of 10 Hz. We model the nominal GPS positioning measurements with the same bounded biases and measurement variances as in the linear double-integrator system. We model the IMU measurement variances using the typical root power spectral density values from an automotive microelectromechanical system (MEMS) inertial module (ST, 2020), and we conservatively model the bounded biases according to the typical range of dynamic bias values for a consumer-grade IMU (Groves, 2013).

We create the simulated trajectory and simulated spoofing trajectory by generating open-loop control input sequences (i.e., yaw rates and longitudinal accelerations) as quartic splines using the technique in Mueller et al. (2015) and then forward-propagating the nonlinear dynamics.

4.2.2 | *Results and Discussion*

For the nominal case of the nonlinear IMU self-contained sensor scenario, similar to the linear propagation model scenario, we observe a CAR that consistently lies above the required rate, as shown in Figure 6(a). Similar to the linear scenario, in the IMU self-contained sensor scenario, the stochastic set of GPS measurement errors $\tilde{\mathcal{P}}_k^{\text{GPS}}$ is constant over all time instances k while the stochastic set of self-contained position errors $\tilde{\mathcal{P}}_k^{\text{self}}$ grows with time, as the IMU errors accumulate with time.

Given that in the nominal case, the biases in the GPS positioning measurements are probabilistically bounded by the corresponding p-zonotope $\mathcal{R}_k$, then the corresponding set $\mathcal{Q}_k$ consistently overbounds the spoofing statistic $\mathbf{q}_k$. As a result, the detector does not raise an alarm, and we qualitatively observe significant intersection of the stochastic sets $\tilde{\mathcal{P}}_k^{\text{GPS}}$ and $\tilde{\mathcal{P}}_k^{\text{self}}$ in the bird's eye view shown in Figure 6(b). In nearly all Monte Carlo trajectories, the detector does not raise an alarm; thus, the Chimera SR-EKF correspondingly uses the fused SR-EKF state estimates and error p-zonotopes for nearly all trajectories, which probabilistically bound the true state, as shown in Figure 6(c).

Similar to the linear propagation model scenario, for the spoofed case of the nonlinear IMU self-contained sensor scenario, we observe that as the spoofing trajectory deviates from the true trajectory, the detector begins to recognize the anomalous GPS measurements and the CDR reaches nearly unity over the Monte Carlo trajectories, as shown in Figure 7(a). Correspondingly, we observe from the bird's eye view of the trajectory in Figure 7(b) that the set of possible states indicated by the GPS position estimate and the self-contained sensor position estimate begin to have little to no overlap, thereby qualitatively validating that the detector should declare a spoofing event.

As the spoofing trajectory deviates from the true trajectory, we observe in Figure 7(b) that the Chimera SR-EKF switches to the self-contained state estimator after detecting a spoofing event. Correspondingly, in Figure 7(c), the Chimera

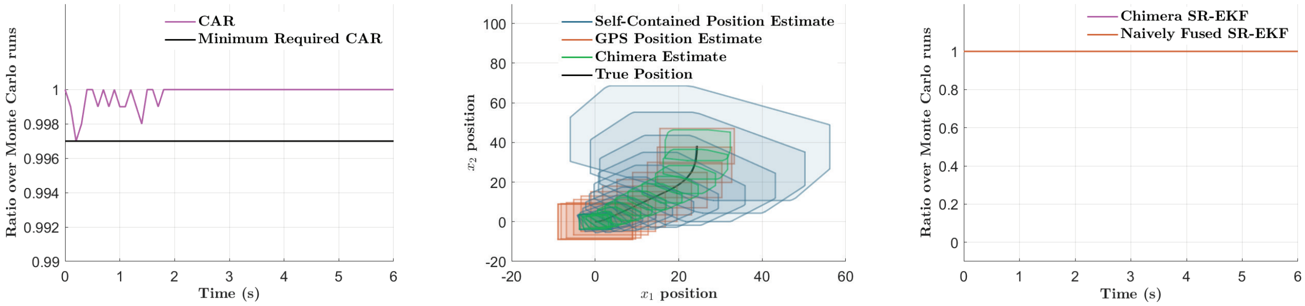

(a) **(b)** **(c)**

FIGURE 6 Monte Carlo validation of the Chimera SR-EKF estimator and detector under nominal conditions for the nonlinear IMU propagation model outlined in Section 4.2.1 (a) Ratio of CAR (b) Bird's eye view of the trajectory (c) Ratio of states bound by 3σ error zonotopes Trajectories and statistics have been averaged over 1000 Monte Carlo simulations of the same trajectory and control. The CAR shown in **(a)** remains above the required value of 0.997, which corresponds to a required FAR of 0.003. In **(b)**, we observe an overlap between the self-contained and GPS 3σ error zonotopes, visually validating our detector's correct authentications on average. For ease of visual interpretation, the error zonotopes are plotted to be centered about the average estimated trajectory across the 1000 Monte Carlo runs, which also coincides with the true trajectory. In **(c)**, we observe that, in the absence of a spoofing attack, the fused SR-EKF estimates bound the true state and are used by the Chimera SR-EKF.

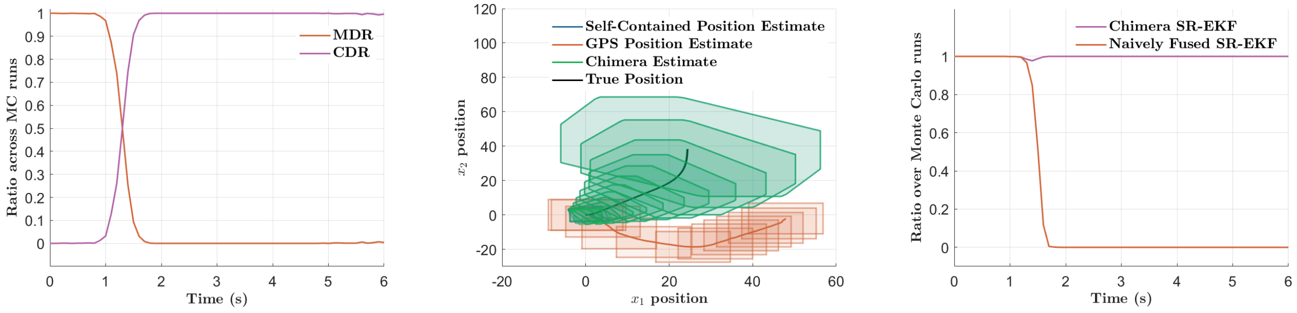

(a) (b) (c)

FIGURE 7 Monte Carlo (MC) validation of the Chimera SR-EKF estimator and detector under spoofed conditions for the nonlinear IMU propagation model outlined in Section 4.2.1 (a) Ratio of MDR and CDR (b) Bird's eye view of the trajectory (c) Ratio of states bound by 3σ error zonotopes

Trajectories and statistics have been averaged over 1000 Monte Carlo simulations of the same trajectory and control. In **(a)**, we observe that the CDR increases, while the MDR correspondingly decreases, as the difference between the two trajectories sufficiently exceeds the errors of the self-contained sensor. For ease of visual interpretation, in **(b)**, the error zonotopes are plotted to be centered about the average Chimera SR-EKF and GPS estimated trajectories across the 1000 Monte Carlo runs, where the average Chimera SR-EKF estimated trajectory also coincides with the true trajectory. We observe that once the trajectories deviate significantly, there is little to no overlap between the self-contained and GPS 3σ error zonotopes. Because our detector recognizes the discrepancy between the sensor measurements and declares a spoofing attack, the Chimera SR-EKF switches to using the self-contained SR filter estimate. Thus, the Chimera SR-EKF error zonotopes are equivalent to those of the self-contained SR-EKF and are correspondingly overlapping for the majority of the trajectory. In **(c)**, we observe that, in the presence of a spoofing attack, the naively fused SR-EKF begins to fail to bound the true state while the Chimera SR-EKF switches to the self-contained state estimate and error zonotope and continues to bound the true state in nearly all Monte Carlo runs.

SR-EKF continues to bound the true state in nearly all Monte Carlo simulations, whereas the naively fused SR-EKF estimate fails to bound the vehicle state once the GPS measurement biases become large.

5 | CONCLUSION

In this work, we derived an SR-based filter and spoofing detector to provide continuously authenticated navigation solutions between Chimera authentication times. Our formal verification method leverages the previously authenticated set of Chimera measurements, in combination with conservative error models for the GPS and self-contained sensor measurements to update the receiver state and uncertainty at each time instant. In particular, we derived an SR detector to satisfy a user-defined false alarm requirement in nominal GPS operation while operating with stochastic errors and unknown bounded biases in the GPS measurements and self-contained sensor measurements. We further extended our state estimation filter and spoofing detector for a nonlinear propagation model by conservatively modeling the linearization error in the state propagation.

While we have focused on the application of our SR filter and detector with the Chimera authentication feature, the techniques and derivations in this work can be applied to any setting in which periodic signal authentication information is available. Through Monte Carlo simulations over a 6-s Chimera fast channel authentication period, we empirically validated for a ground vehicle model that our Chimera SR-KF and SR-EKF detectors satisfy the user-defined false alarm requirement during the Chimera epoch, while detecting spoofing during simulated trajectory-drifting spoofing attacks. Additionally, we demonstrated that our Chimera SR-KF and SR-EKF estimators successfully bound the vehicle state under both authentic and spoofing conditions.

ACKNOWLEDGMENTS

We thank Akshay Shetty for insightful discussions during the development of this work and for reviewing drafts of this paper. We also thank Shubh Gupta for reviewing drafts of this paper.

This material is based upon work supported by the AFRL at Kirtland Air Force Base in Albuquerque, New Mexico, under grant number FA9453-20-1-0002. This material is additionally based upon work supported by a National Science Foundation Graduate Research Fellowship under grant number DGE-1656518. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the AFRL or the National Science Foundation.

REFERENCES

AFRL. (2023). Navigation Technology Satellite - 3 NTS-3. Air Force Research Laboratory. https://afresearchlab.com/technology/nts-3

AFRL Space Vehicles Directorate, Advanced GPS Technology. (2019). Chips message robust authentication (Chimera) enhancement for the L1C signal: Space segment/user segment interface. *IS-AGT-100*. https://drive.google.com/file/d/1r_3F7q4aItLt3ANYd847EoN3NxFKRNb0/view

Althoff, M., Grebenyuk, D., & Kochdumper, N. (2018). Implementation of Taylor models in CORA 2018. *Proc. of the 5th International Workshop on Applied Verification for Continuous and Hybrid Systems (ARCH18),* Oxford, United Kingdom, 145–173. https://doi.org/10.29007/zzc7

Althoff, M. (2010). *Reachability analysis and its application to the safety assessment of autonomous cars* [Doctoral dissertation, Technische Universität München]. https://mediatum.ub.tum.de/doc/1287517

Althoff, M., Stursberg, O., & Buss, M. (2008). Reachability analysis of nonlinear systems with uncertain parameters using conservative linearization. *Proc. of the 47th IEEE Conference on Decision and Control*, Cancun, Mexico, 4042–4048. https://doi.org/10.1109/CDC.2008.4738704

Althoff, M., Stursberg, O., & Buss, M. (2009). Safety assessment for stochastic linear systems using enclosing hulls of probability density functions. *Proc. of the 2009 European Control Conference (ECC)*, Budapest, Hungary, 625–630. https://doi.org/10.23919/ECC.2009.7074473

Anderson, J. M., Carroll, K. L., DeVilbiss, N. P., Gillis, J. T., Hinks, J. C., O'Hanlon, B.W., Rushanan, J. J., Scott, L., & Yazdi, R. A. (2017). Chips-message robust authentication (Chimera) for GPS civilian signals. *Proc. of the 30th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2017)*, Portland, OR, 2388–2416. https://doi.org/10.33012/2017.15206

Berz, M., & Hoffstätter, G. (1998). Computation and application of Taylor polynomials with interval remainder bounds. *Reliable Computing*, *4*(1), 83–97. https://doi.org/10.1023/A:1009958918582

Bhamidipati, S., & Gao, G. X. (2020a). GPS spoofing mitigation and timing risk analysis in networked PMUs via stochastic reachability. *Proc. of the 33rd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2020)*, 3920–3937. https://doi.org/10.33012/2020.17757

Bhamidipati, S., & Gao, G. X. (2020b). Integrity-driven landmark attention for GPS-vision navigation via stochastic reachability. *Proc. of the 33rd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2020)*, 2311–2326. https://doi.org/10.33012/2020.17546

Blanch, J., Walter, T., & Enge, P. (2017). A MATLAB toolset to determine strict Gaussian bounding distributions of a sample distribution. *Proc. of the 30th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2017)*, Portland, OR, 4236–4247. https://doi.org/10.33012/2017.15392

Blanch, J., Walter, T., & Enge, P. (2018). Gaussian bounds of sample distributions for integrity analysis. *IEEE Transactions on Aerospace and Electronic Systems*, *55*(4), 1806–1815. https://doi.org/10.1109/TAES.2018.2876583

Broumandan, A., & Lachapelle, G. (2018). Spoofing detection using GNSS/INS/odometer coupling for vehicular navigation. *Sensors*, *18*(5), 1305. https://doi.org/10.3390/s18051305

Combastel, C., & Zolghadri, A. (2020). A distributed Kalman filter with symbolic zonotopes and unique symbols provider for robust state estimation in CPS. *International Journal of Control*, *93*(11), 2596–2612. https://doi.org/10.1080/00207179.2019.1707278

Cozzens, T. (2021). NTS-3 mission progresses toward launch in 2023. *GPS World*. https://www.gpsworld.com/nts-3-mission-progresses-toward-launch-in-2023/

DeCleene, B. (2000). Defining pseudorange integrity-overbounding. *Proc. of the 13th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS 2000)*, Salt Lake City, UT, 1916–1924. https://www.ion.org/publications/abstract.cfm?articleID=1603

Divis, D. A. (2019). New Chimera signal enhancement could spoof-proof GPS receivers. *Inside GNSS*. https://insidegnss.com/new-chimera-signal-enhancement-could-spoof-proof-gps-receivers/

GPS Directorate. (2022). *NAVSTAR GPS space segment/user segment L1C interfaces* (tech. rep.). IS-GPS-800J. https://www.gps.gov/technical/icwg/IS-GPS-800J.pdf

Groves, P. D. (2013). *Principles of GNSS, inertial, and multisensor integrated navigation systems*. 2nd ed. Artech House.

Humphreys, T. E. (2013). Detection strategy for cryptographic GNSS anti-spoofing. *IEEE Transactions on Aerospace and Electronic Systems*, *49*(2), 1073–1090. https://doi.org/10.1109/TAES.2013.6494400

Kousik, S., Dai, A., & Gao, G. (2021). Ellipsotopes: Combining ellipsoids and zonotopes for reachability analysis and fault detection. *arXiv preprint arXiv:2108.01750*. https://doi.org/10.48550/arXiv.2108.01750

Kousik, S., Holmes, P., & Vasudevan, R. (2019). Safe, aggressive quadrotor flight via reachability-based trajectory design. *Proc. of the ASME 2019 Dynamic Systems and Control Conference Vol. 59162*, Park City, Utah, V003T19A010. https://doi.org/10.1115/DSCC2019-9214

Medina Lee, J. F., Trentin, V., & Villagra, J. (2019). Framework for motion prediction of vehicles in a simulation environment. *XL Jornadas de Automática*, 520–527. https://doi.org/10.17979/spudc.9788497497169.520

Mina, T., Kanhere, A., Kousik, S., & Gao, G. (2021). Continuous GPS authentication with Chimera using stochastic reachability analysis. *Proc. of the 34th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2021)*, St. Louis, MO, 4234–4248. https://doi.org/10.33012/2021.18125

Mueller, M. W., Hehn, M., & D'Andrea, R. (2015). A computationally efficient motion primitive for quadrocopter trajectory generation. *IEEE Transactions on Robotics*, *31*(6), 1294–1310. https://doi.org/10.1109/TRO.2015.2479878

Rife, J., Pullen, S., Enge, P., & Pervan, B. (2006). Paired overbounding for nonideal LAAS and WAAS error distributions. *IEEE Transactions on Aerospace and Electronic Systems*, *42*(4), 1386–1395. https://doi.org/10.1109/TAES.2006.314579

Rife, J., Pullen, S., Pervan, B., & Enge, P. (2004). Paired overbounding and application to GPS augmentation. *Proc of the IEEE Position, Location, and Navigation Symposium (PLANS 2004)*, Monterey, CA, 439–446. https://doi.org/10.1109/PLANS.2004.1309027

Schürmann, B., Klischat, M., Kochdumper, N., & Althoff, M. (2021). Formal safety net control using backward reachability analysis. *IEEE Transactions on Automatic Control*, *67*(11), 5698–5713. https://doi.org/10.1109/TAC.2021.3124188

Shetty, A., & Gao, G. X. (2019). Predicting state uncertainty for GNSS-based UAV path planning using stochastic reachability. *Proc. of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2019)*, Miami, FL, 131–139. https://doi.org/10.33012/2019.16896

Shetty, A., & Gao, G. X. (2021). Predicting state uncertainty bounds using non-linear stochastic reachability analysis for urban GNSS-based UAS navigation. *IEEE Transactions on Intelligent Transportation Systems*, *22*(9), 5952–5961. https://doi.org/10.1109/TITS.2020.3040517

ST. (2020). *Automotive 6-axis inertial module: 3D accelerometer and 3D gyroscope* [Rev. 4]. https://www.st.com/en/mems-and-sensors/asm330lhh.html

Tanıl, C., Khanafseh, S., Joerger, M., & Pervan, B. (2017). An INS monitor to detect GNSS spoofers capable of tracking vehicle position. *IEEE Transactions on Aerospace and Electronic Systems*, *54*(1), 131–143. https://doi.org/10.1109/TAES.2017.2739924

Wesson, K., Rothlisberger, M., & Humphreys, T. (2012). Practical cryptographic civil GPS signal authentication. *NAVIGATION, 59*(3), 177–193. https://doi.org/10.1002/NAVI.14

How to cite this article: Mina, T., Kanhere, A., Kousik, S., & Gao, G. (2023). Stochastic reachability-based GPS spoofing detection with Chimera signal enhancement. *NAVIGATION, 70*(4). https://doi.org/10.33012/navi.616