Assessing GNSS Spoofing Impact on A Safety-Critical Land Transportation Localization Function Within a Cooperative Fleet: An End-Users Focused Experimental Study

Zaynab El Mawas, Nourdine Ait Tmazirte, Cindy Cappelle, Maan El Badaoui El Najjar

Assessing GNSS Spoofing Impact on A Safety-Critical Land Transportation Localization Function Within a Cooperative Fleet: An End-Users Focused Experimental Study

Zaynab EL MAWAS[1], Nourdine AIT TMAZIRTE[2], Cindy CAPPELLE[1], Maan EL BADAOUI EL NAJJAR[1]

[1]*Univ. Lille, CNRS, Centrale Lille, UMR 9189 CRIStAL, F-59000 Lille, France*

[2]*Univ Gustave Eiffel, COSYS-LEOST, F-59650 Villeneuve d'Ascq, France*

email: {zaynab.elmawas, cindy.cappelle, maan.el-badaoui-el-najjar}@univ-lille.fr, {nourdine.ait-tmazirte}@univ-eiffel.fr

ABSTRACT

The widespread adoption of Global Navigation Satellite Systems (GNSS) in land transportation highlights the critical security concerns posed by GNSS spoofing. These threats compromise the reliability and safety of transportation, impacting not only individual vehicles but also the collective functionality of these systems. Prior studies have primarily concentrated on the detection and mitigation of GNSS spoofing at the level of individual vehicles. This research breaks new ground by examining the effects of GNSS spoofing on a fleet of vehicles, an aspect that has not been explored in prior work. By leveraging a combination of advanced GNSS signal emulation platforms like Safran Skydel SDX and a practical setup involving well-equipped robotized cars, the dynamics of how spoofing influences fleet operations under various conditions are examined. The study assesses the degradation in fleet performance through a series of controlled experiments. The investigation highlights how spoofing impacts Key Performance Indicators (KPIs) such as accuracy, availability and integrity within cooperative fleets, offering insights into the development of robust protective strategies to ensure the integrity and safety of future cooperative land transportation systems.

Keywords: Resilient GNSS, Spoofing of GNSS signals, Cooperative localization, Integrity monitoring.

I. INTRODUCTION

The growing use of Global Navigation Satellite Systems (GNSS) in safety-critical land transportation systems, such as autonomous vehicles, has led to an increased focus on ensuring the accuracy, availability, and safety of GNSS-based positioning systems. GNSS technology enables a fast and cost-effective means of accurately determining one's position on Earth (and for some applications even beyond) without prior knowledge, making it an essential component of modern transportation systems. However, this technology is inherently vulnerable to various natural threats, including atmospheric delays, imprecise satellite ephemeris, building reflections resulting in multipath or NLOS, scattering effects of trees, and artificial threats such as jamming, and spoofing. Spoofing is a malicious attack where a fake GNSS signal is generated to deceive a GNSS receiver into tracking the fake signal instead of the genuine one. It is defined in automotive standards (ISO/SAE 21434) as an intentional and unauthorized transmission of GNSS-like signals with the purpose of deceiving a GNSS receiver. In safety-critical transportation systems, spoofing can result in significant degradation of positioning accuracy, reduced availability, and safety concerns.

Based on a block diagram of a receiver, multiple strategies have been introduced for the detection of GNSS spoofing, varying in their levels of implementation.

At the satellite level, methods that utilize cryptographic techniques for authentication serve to confirm the integrity of GNSS signals. For instance, the Galileo system offers the Open Service Navigation Message Authentication (OSNMA) within its E1-B signal component. This involves a cryptographic protocol that delivers unpredictable data symbols to users, enabling them to validate the I/NAV message's content. In this context, Shahid et al. [2023] introduced a detection method for spoofers that analyzes snapshots of received unpredictable symbols and juxtaposes them with genuine symbols. Furthermore, Humphreys [2013] introduced a detection strategy for spoofing attacks on both military and enhanced-civil cryptographically-secured GNSS signals, employing a hypothesis test tailored to replay-type attacks.

At the antenna level, the implementation of directional antennas supports the identification of incoming signal directions and the pinpointing of spoofing sources, as shown by Kang et al. [2018] through the comparison of estimated direction-of-arrival (DOA) data with actual DOA information sourced from GPS almanacs and ephemerides. Psiaki et al. [2013] proposed a method for detecting spoofing in GPS signals using a dual-receiver correlation approach, involving a reference receiver and a defended

receiver that could be the target of a spoofing attack.

At the receiver stage, focus is given to signal processing methods like signal fingerprinting, improved signal matching, and checking signal strength. Akos [2012] used automatic gain control (AGC) in the radio frequency section to identify GPS/GNSS spoofing.

At a broader system level, additional sensors such as Inertial Measurement Units (IMUs), radar, lidar, or cameras are used. Broumandan and Lachapelle [2018] developed a method that checks the consistency and analyzes GNSS and IMU/odometer data separately within a specific time frame, then compares these results with those from GNSS and the inertial navigation system (INS)/odometer.

Regarding infrastructure-based solutions, GNSS augmentation systems are particularly interesting for checking signal integrity more thoroughly and for using wireless networks to aid in verifying GNSS positioning. Additionally, approaches that use data analysis, pattern recognition, and machine learning review past GNSS signal patterns and receiver actions to spot unusual behavior that might suggest spoofing. Shafiee et al. [2017] introduced a way to detect GPS spoofing using a multi-layer neural network that distinguishes between fake and real signals by looking at differences in phase, energy, and the imaginary parts of the signals.

Various methods have been proposed to categorize spoofing attacks based on their application domains. Fernendez-Hernendez et al. [2019] classify interference into four levels of jamming and seven levels of spoofing, considering the intention and sophistication of the attacks. Other researchers, such as Merwe et al. [2018], Dasgupta et al. [2022], distinguish between synchronous attacks, which transmit spoofing signals with overlapping correlation peaks, and asynchronous attacks, which transmit non-time-synchronized signals. Each underlies multiple sub-classes given the application of the spoofing. In the synchronous attacks, Rothmaier [2021] categorizes synchronous spoofing into two types: jam-then-spoof attacks and lift-off attacks. In a jam-then-spoof attack, the attacker initially suppresses the authentic signal by broadcasting noise on the carrier frequency of interest, preventing the receiver from decoding the GNSS signals. This allows the spoofer's signals to dominate, causing the receiver to lock onto them. In lift-off attacks, the spoofer carefully aligns the malicious signals with the authentic ones before moving the correlation peak and altering the resulting pseudoranges as desired.

The literature indicates that addressing malicious spoofing necessitates a multi-level approach, wherein each layer operates synergistically with the others to strengthen overall effectiveness.

While there are established techniques for detecting and mitigating GNSS spoofing in a single vehicle or receiver, far less attention has been given to the cooperative case, where multiple vehicles share positioning data. In cooperative and safety-critical systems like vehicle fleets, spoofing poses an even greater risk, as a compromised vehicle can propagate errors throughout the entire fleet, affecting both safety and efficiency.

This study aims to address the existing gap by evaluating the impact of spoofing on the various Stanford diagrams of each vehicle within the fleet. The primary goal is to establish the experimental framework necessary to assess, in the first phase, the impact of an unknowingly spoofed vehicle taking part in a cooperative system —in both directions: the negative effect of the spoofed vehicle on the rest of the fleet and the positive effect of the fleet on the spoofed vehicle. While this study will not implement any detection or mitigation techniques, it will focus on providing an in-depth and detailed analysis based on a realistic scenario.

II. RESEARCH PROTOCOL

In order to address the problem of GNSS spoofing in cooperative vehicle fleets, we have defined a detailed experimental protocol in Fig 1.

The first step is the real data acquisition, which involves using a fleet of three robotized Renault Zoé vehicles from the PRETIL platform of the CRIStAL laboratory. Next, we move to the controlled environment for spoofing, where we simulate GNSS spoofing using Safran's Skydel GNSS simulator. The spoofing attack available in the Skydel simulation is the jam-then-spoof attack, classified as S4 Targeted Spoofers – Simulation given by the categorisation of Fernendez-Hernendez et al. [2019]. This allows us to generate spoofed positions for one of the vehicles while maintaining real, unspoofed data for the others. Finally, we proceed to the positioning and performance analysis phase. Here, we apply the multi-sensor fusion approach to compute protection levels in various scenarios, comparing the results from the spoofed vehicle versus non-spoofed, as well as cooperative versus non-cooperative cases. This analysis is done using Stanford diagrams and statistical methods to evaluate how spoofing impacts the overall fleet's accuracy, as well as the potential corrective effects of cooperation within the fleet.

1. Data Acquisition

For the real-world data acquisition, all vehicles were managed through the ROS (Robotic Operating System) middleware for sensor data acquisition and synchronization. The vehicles were equipped with GNSS receivers, Inertial Navigation Systems (INS), LiDARs, and CAN odometry systems, which were loosely coupled for integration (See Fig 2).

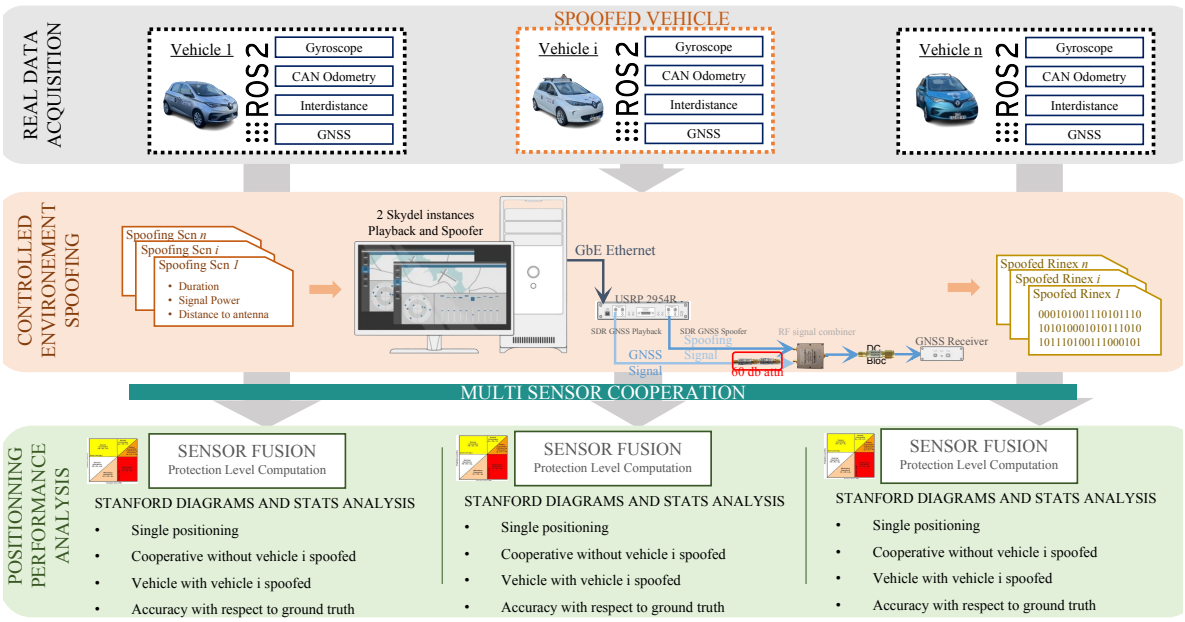

Figure 1: Experimental Protocol

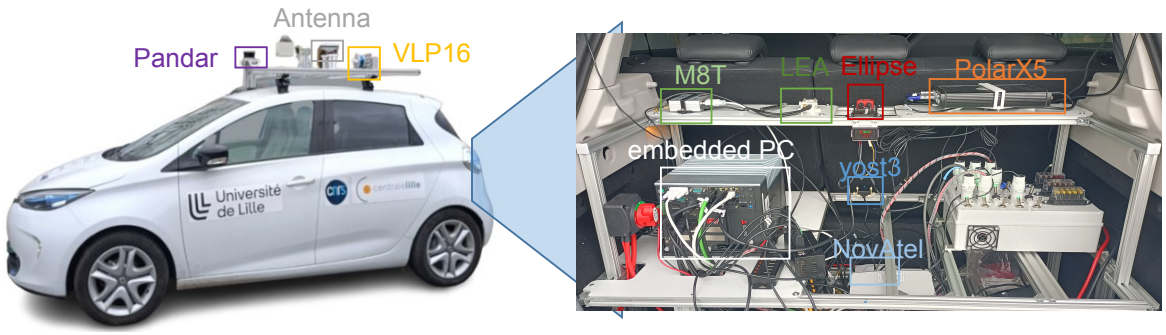

Figure 2: Equipped Sensors

The sensors were strategically positioned in the vehicle's trunk, establishing distinct reference frames for each. Specifically, the IMUs were aligned with the axis of the rear wheels, which serves as the vehicle's base frame for movement. This alignment ensures accurate motion and position tracking relative to the vehicle's dynamics.

For ground truth regarding position and orientation, a Novatel receiver was employed, with data collected and synchronized through ROS. In contrast, the Septentrio PolaRX5 receiver was managed via its proprietary software, rather than ROS, to generate navigation data. Additional position data was obtained from the Ublox EVK-M8t receiver, with ROS facilitating the data acquisition. To ensure consistent signal reception, a shared antenna was used for all receivers, enabling uniform GNSS data across devices.

Alongside the GNSS data, angular velocity was captured using both SBG Ellipse and YOST3 IMUs. Additionally, wheel encoder data was gathered from the ABS system and CAN bus. This odometric data was incorporated into a model for predicting vehicle movement, contributing to the overall data fusion process.

The data acquisition process began with all vehicles positioned at a common starting location. Following this initial phase, the vehicles were directed towards different trajectories and locations. To preserve the integrity of the experiment, the vehicle paths were carefully adjusted to avoid overlap, ensuring that spoofing directly affected only one vehicle (Car1) without influencing others. The dataset for this study is publicly available and can be accessed through Dherbomez et al..

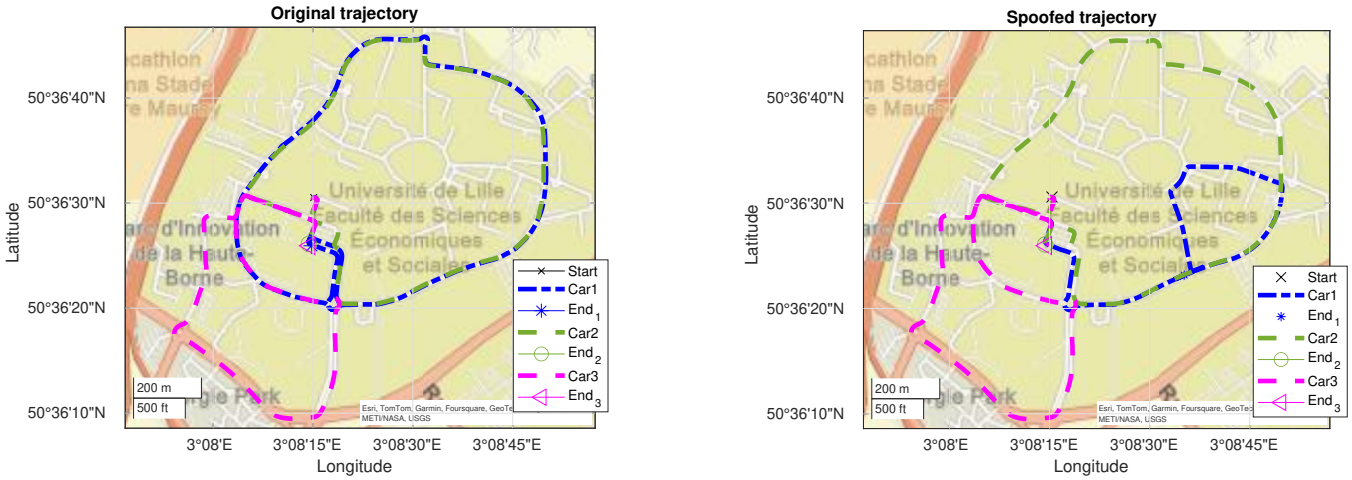

Figure 3: Original acquired trajectory and spoofed recorded trajectory

2. Spoofing Simulation

For the controlled spoofing simulation, we employ Safran's Skydel SDX software, running two distinct instances. The first instance plays back the original position data based on the vehicle's speed profile, while the second instance is responsible for the spoofing operation, where the position of the vehicle is deliberately manipulated to deviate from its true path.

In our spoofing scenario, the spoofing area was carefully designed with a 250-meter radius (500-meter diameter). The spoofed trajectory was simulated with the intention of rerouting the vehicle in a loop, preventing any intersection with other vehicles. This ensures that the spoofed vehicle does not encounter any other vehicle that could potentially provide corrective data via perception sensors (such as LiDAR or cameras) and communication.

The spoofing simulation is controlled by the several key parameters:

- **Constellation** : GPS L1 C/A and SBAS L1
- **Timing of Signal Activation**:
  - **Jammer Activation**: The jammer is activated for 5 to 10 seconds, long enough for the GNSS receiver to lose correlation with the legitimate signal.
  - **Spoofer Activation**: The spoofing signal is activated 4 seconds after the jammer starts. Once the spoofed signal is locked, the jammer is deactivated after an additional 4 seconds, allowing the spoofed signal to dominate.
- **Jammer Signal Properties**:
  - The jammer uses a chirp signal that covers the entire GPS L1 C/A bandwidth (2.046 MHz), with a short sweep time of 100 µs to prevent the GNSS receiver from locking onto it.

The spoofing process unfolds as follows:

- **Initial Playback**: The original trajectory, based on the vehicle's speed profile, is played on the main instance, and the altered one is played on the spoofing instance.
- **Jammer Activation**: Upon entering the spoofing area, the jammer is activated at 02min18sec, disrupting the vehicle's reception of legitimate GNSS signals.
- **Spoofer Activation**: Four seconds after the jammer starts, the spoofing signal is activated, and 4 seconds later, the jammer is deactivated, allowing the vehicle to lock onto the spoofed signal.
- **Receiver Response**: After jamming and spoofing, the covariance of the receiver's position increases, indicating rising uncertainty. Within 3 seconds of jamming deactivation, the vehicle locks onto the spoofed signal and follows the spoofed trajectory, diverging from the original path.

To ensure the spoofing works as intended, the vehicle's speed is restricted to 50 km/h within the spoofed area of 500 meters. This allows the receiver to lose correlation with the legitimate GNSS signal within the first 30% of the trajectory (150 meters).

The trajectories of the vehicles, along with the position of the jammer, spoofer and the spoofing zone are presented in Figure 4, where the blue cross indicates the position of start of spoofing of Car 1.

It is important to note that none of the other vehicles in the fleet were inside the jamming area when the jamming and spoofing took place. Therefore, these vehicles remained locked onto the authentic GNSS signal throughout the experiment. The unaffected vehicles maintained their proper positions, continuing to operate based on valid GNSS data, without any influence from the spoofing event.

This design ensured that the spoofed vehicle's erroneous behavior could be observed independently, while the rest of the fleet remained unspoofed. The separation of spoofed and non-spoofed vehicles allowed us to study how cooperative localization—sharing position data from non-spoofed vehicles—might correct the spoofed vehicle's position and how this erroneous data impacted the fleet's overall behavior.

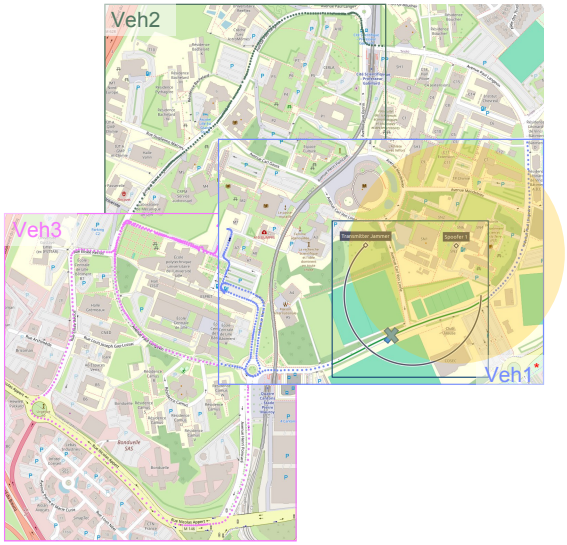

Figure 4: Trajectories of the cars and the location of the spoofer and the jammer

For the controlled spoofing simulation, we chose the Ublox EVK-M8T receiver after testing various receivers, including Septentrio PolarX and Novatel pwrpak, under the nominal Skydel spoofing scenario as well as our case study conditions. The decision was based on three key factors: its large availability, its vulnerability to the discussed spoofing issue, and its adaptability with our experimental framework, particularly the integration with ROS for data registration. This made the Ublox receiver an ideal choice for our study. However, it is important to note that the M8T series is not the latest model, and our findings do not provide insights into the resilience of newer models in the series.

III. POSITION ESTIMATION AND PERFORMANCE MONITORING

In this study, we employ a stochastic, non-deterministic approach to position estimation, focusing not only on the estimated position but also on the covariance and its evolution to observe the system's integrity. This is achieved through the use of Bayesian filters, which allow us to account for uncertainties in the measurements and provide probabilistic estimates of the vehicle's state.

The vehicle is modeled with three degrees of freedom (3-DOF), encompassing both its position and orientation. At any given time instant k, the state vector is defined as the vehicle's position in the ENU (East, North, Up) reference frame, along with its heading angle, θ. The state vector is represented as:

$$X_k = [x_{ENU} \quad y_{ENU} \quad \theta_{ENU}]^T \tag{1}$$

The choice of ENU frame was due to its better alignment with the movement of the vehicle on a planar surface, simplifying the integration of odometric and sensor data into the position estimation process. Additionally, the ENU frame avoids the complexity of managing curved earth effects that are more prominent in ECEF for larger-scale applications.

1. Sensor fusion and position estimation

For position estimation, we utilize the Kalman Filter in its Extended Information Filter (EIF) form, which enables effective multi-sensor fusion. This process involves two key steps: prediction and correction.

In the prediction step, we rely on the odometric model derived from the rear wheel encoders. This model helps estimate the vehicle's position based on its motion over time.

The correction step involves summing the informational contributions from various observations, applying therefore sensor fusion approach. In this case study, we use three main correction sources:

- GNSS correction (Z_{GNSS}): which is obtained at the antenna level and then transformed to the vehicle's base frame for alignment with the vehicle's coordinate system.
- IMU correction (Z_{gyro}): the gyroscope's angular velocity data is utilized to correct the orientation of the vehicle. The gyroscope provides continuous measurements of the vehicle's angular velocity, which are integrated to estimate the change in orientation over time.
- Relative correction ($Z_{rel}^{i \to j}$): in this cooperative localization system and to ensure that only one vehicle is spoofed, vehicles are positioned far apart, making the use of typical perception sensors like LiDAR or cameras infeasible. Instead, we simulate inter-vehicle perception using ground truth information, inspired by the use of the Ultra-Wide-Band (UWB) technology for cooperative localization as described in Pierre et al. [2018]. This allows us to mimic UWB-based measurements for both distance and bearing between vehicles.

  Building on our previous work Mawas et al. [2023], where Range Finders were used to estimate inter-vehicle distances and bearing angles, the vehicles act as a mobile beacon. They share their position estimation over a closed communication network, allowing each vehicle to localize itself based on the positions of the other vehicles in the fleet and its relative position to them (see Figure 5).

  For a fleet of N vehicles, each vehicle can make use of up to $N-1$ relative observations to refine its position estimate. A relative observation from Vehicle 1 to Vehicle 2 can be obtained as the following :

$$Z_{rel}^{1\to2} = \begin{bmatrix} Z_x^{rel^{1\to2}} \\ Z_y^{rel^{1\to2}} \end{bmatrix} = \begin{bmatrix} x^2 \\ y^2 \end{bmatrix} - \begin{bmatrix} \rho^{1\to2} \times \cos(Z_b^{1\to2} + \theta^1) \\ \rho^{1\to2} \times \sin(Z_b^{1\to2} + \theta^1) \end{bmatrix} \tag{2}$$

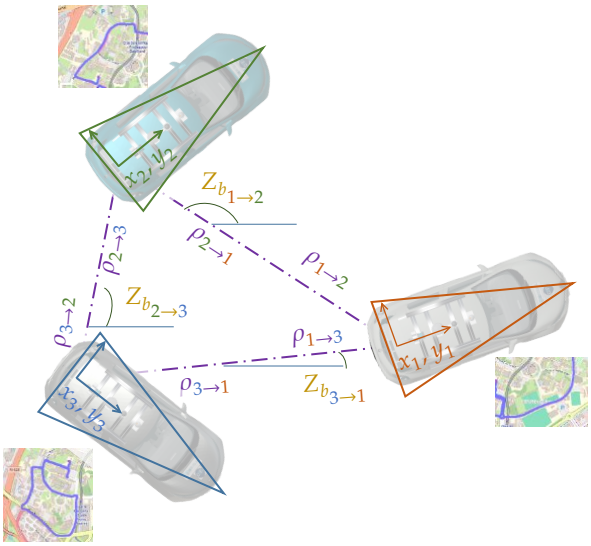

Figure 5: Simulation of the relative observation between the vehicles

In multi-sensor data fusion, the information matrix from the prediction phase $Y_{k|k-1}$ is augmented by the combined informational contributions from the sensors $I_k^l = (H_k^l)^T \times (R_k^l)^{-1} \times H_k^l$, using the following equations within the information filter

framework:

$$\begin{cases} Y_{k|k} = Y_{k|k-1} + \sum_{l=1}^{L} (H_k^l)^T \times (R_k^l)^{-1} \times H_k^l \\ y_{k|k} = y_{k|k-1} + \sum_{l=1}^{L} (H_k^l)^T \times (R_k^l)^{-1} \times Z^{obs^l} \end{cases} \tag{3}$$

Where in our case, having $L = 4$ observations, therefore $Z^{obs} = [Z_{GNSS}, Z_{gyro}, Z_{rel}^{1\to2}, Z_{rel}^{1\to3}]$.

2. Performance monitoring

In this study on the impact of spoofing, the key performance indicators (KPIs) are the accuracy, the availability and the potential misleading information (linked to the safety) of the estimated position. For that, we utilize a Stanford diagram, focusing specifically on the horizontal component of the position covariance.

To evaluate integrity, the Horizontal Protection Level (HPL) is calculated as six times the square root of the largest eigenvalue of the covariance matrix, based on the method from Boysen and Zunker [2005]. This ensures a 99.9999% probability of accuracy, which is necessary for high-integrity applications such as cooperative navigation and spoofing impact assessments. The alert limit (HAL) is set to 3.5 meters, which corresponds to half the average width of a two-lane road (see Figure 6).

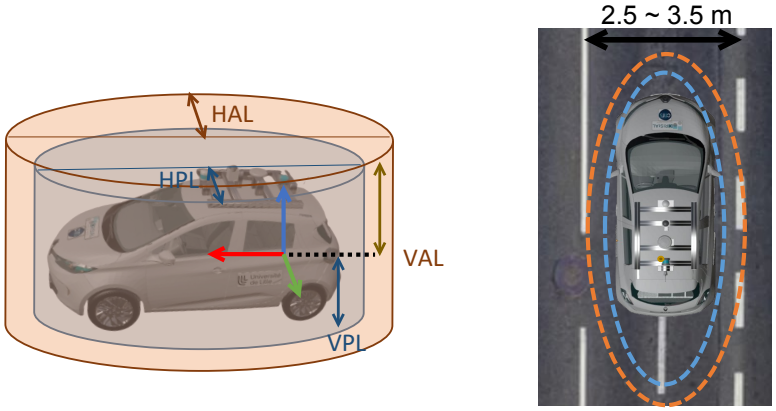

Figure 6: Protection levels

The formula for the HPL is expressed as:

$$HPL = 6\sqrt{\lambda_{(max)}} = 6\sqrt{(max(eigen(P_{k|k})))}. \tag{4}$$

Here, $\lambda_{(max)}$ represents the largest eigenvalue of the position covariance matrix $P_{k|k}$ which is evaluated at time step k. This high factor of 6 ensures that the system's integrity risk is sufficiently minimized, aligning with the integrity bounds commonly applied in safety-critical systems.

We emphasize the use of the HPL because it captures the overall horizontal uncertainty in the vehicle's estimated position, which is critical for maintaining lateral accuracy on the road. Importantly, this concept can be further expanded to analyze both Along-Track (AT) and Cross-Track (CT) errors in future work. By breaking down the HPL into two components—representing position uncertainty along the vehicle's direction of travel (AT) and across its lateral movement (CT)—we can more accurately assess how spoofing or other errors affect different aspects of vehicle positioning.

IV. RESULTS AND DISCUSSION

In this section, we present the results from our study, focusing on two key localization strategies: single localization and cooperative localization. The goal of this work is to evaluate the effects of spoofing on both spoofed and non-spoofed vehicles within these frameworks, while contributing insights into how cooperative systems can influence localization accuracy and integrity under attack.

The study considers four main scenarios:

1. Single Localization Non-Spoofed (SL-NS): In this case, each vehicle localizes itself independently without any external influence. This serves as the baseline for evaluating the performance of the system in normal conditions.
2. Single Localization Spoofed (SL-S): Here, one vehicle is subject to spoofing, where false GPS data is injected into the system. We analyze how this impacts the vehicle's ability to maintain an accurate position estimate when relying solely on its own sensors and positioning data.
3. Cooperative Localization Non-Spoofed (CL-NS): In this scenario, vehicles share position data with each other through a cooperative framework. This case illustrates the benefit of cooperative localization in improving accuracy and integrity when no spoofing is involved.
4. Cooperative Localization Spoofed (CL-S): This is the critical scenario in which one vehicle is spoofed, but position data is still shared among the fleet. We evaluate both the negative impact of the spoofed vehicle on the rest of the fleet, as well as how the shared data from unaffected vehicles can positively influence the spoofed vehicle's localization.

1. Single Localization No spoofing case (SL-NS)

In this first scenario, we analyze the performance of the localization system in the Single Localization mode, with no spoofing applied to any vehicle. The results are presented in Figure 7. Notably, there are no occurrences of Hazardous Misleading Information (HMI) or System Unavailability across all vehicles, indicating that the system operates within safe and reliable bounds. The majority of the points for all vehicles are found in the Normal Operation region, with the average localization error being around 2.3 meters.

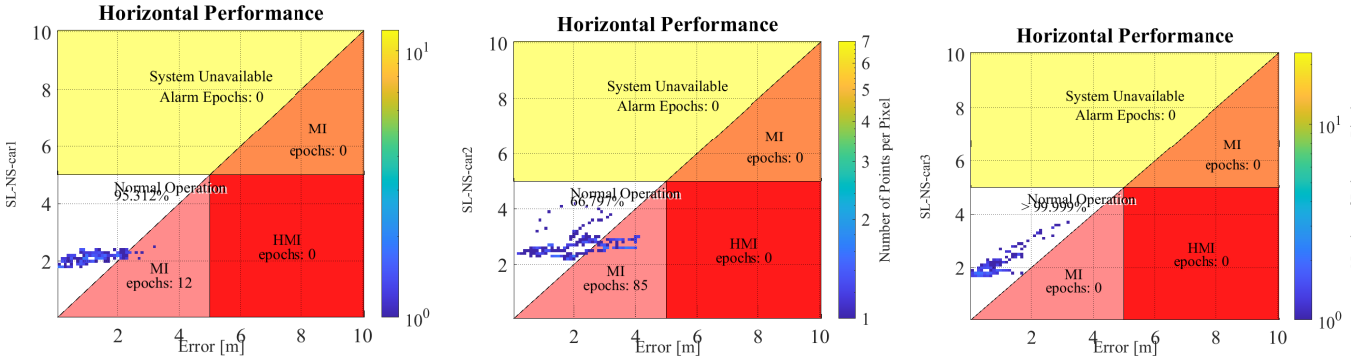

(a) Car 1 **(b)** Car 2 **(c)** Car 3

Figure 7: Integrity monitoring for the case SL-NS for the 3 vehicles

Car 1 and Car 3 show similar strong performance, with most of their data points concentrated within the Normal Operation zone. Their estimated positions remain close to the ground truth, and the system successfully maintains high precision and integrity. Car 2, however, demonstrates reduced performance compared to the other vehicles. While most of its points remain in the Normal Operation region, a significant number drift closer to the HMI zone, indicating that its position estimates are less precise and occasionally approach unsafe error levels.

The performance discrepancy between Car 2 and Car 3 is noteworthy and beneficial for the next phases of the study, where we analyze the impact of cooperation and spoofing. The fact that Car 2 consistently shows lower performance, with points approaching the HMI region, provides a valuable contrast to the stronger performance of Car 3. This diversity allows us to assess how cooperation between vehicles with varying localization accuracy can influence the overall system.

2. Single Localization Spoofing (SL-S)

After applying spoofing to Vehicle 1, we observe significant degradation in its localization performance, as shown in the Stanford diagram. While Cars 2 and 3 remain unaffected due to the single localization setup, Car 1 shows a significant shift in its performance (see Figure 8).

Before the spoofing attack, Car 1 performs similarly to the no-spoofing case, with approximately 95% of the points residing in the Normal Operation zone. However, as spoofing takes effect, only 35% of the points remain in the Normal Operation region—those that occurred prior to or at the very beginning of the spoofing event.

At the start of the test, before the spoofing is applied, Car 1 maintains a localization error around 2 to 2.5 meters, which is consistent with its performance in the no-spoofing scenario. This precision is acceptable and falls within the system's integrity

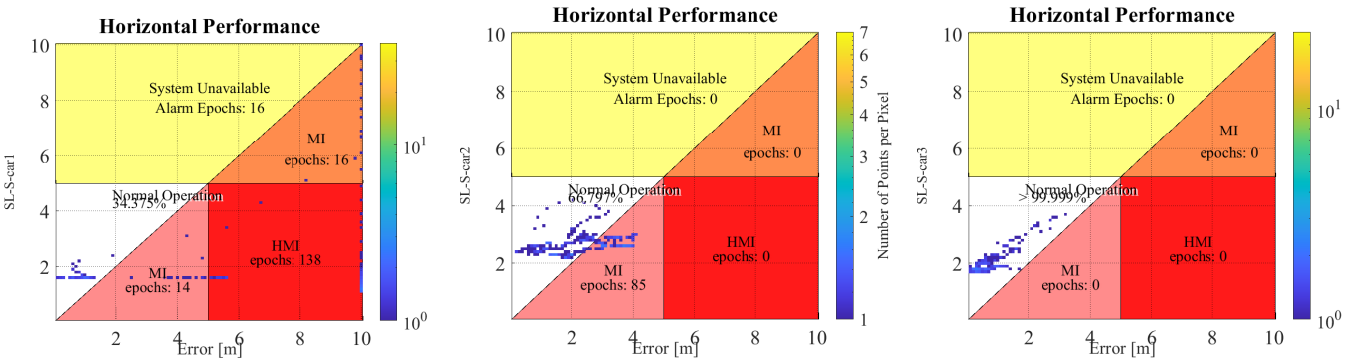

(a) Car 1 **(b)** Car 2 **(c)** Car 3

Figure 8: Integrity monitoring for the case SL-S

bounds.

As the spoofing takes hold, the error begins to rise gradually, pushing more points into the MI region. Here, the localization error becomes significant, exceeding acceptable thresholds without the system triggering an alert. This shift is critical because it means the vehicle is operating under incorrect position estimates without any indication that its localization is deteriorating.

Over time, as the spoofing attack continues, Car 1's position error increases further, and several points transition into the HMI region. This signifies that the vehicle's localization error is not only large but also undetected, posing a significant safety risk. If the vehicle continues to rely on these erroneous position estimates, it could result in dangerous driving decisions, such as incorrect lane-keeping or path planning, which is the case with spoofing.

Since this is a Single Localization scenario, the spoofing attack is isolated to Car 1. As a result, Cars 2 and 3 continue to maintain their previous performance levels, with no changes in their localization accuracy. This illustrates the isolated nature of spoofing in a single localization system, where the error does not propagate to other vehicles.

3. Cooperative Localization No spoofing (CL-NS)

In this scenario, we implement the Cooperative Localization approach, where each vehicle shares its location data over a closed network, and all vehicles use the positions of others as mobile landmarks.

The introduction of cooperative localization has a noticeable positive impact on the performance of all the vehicles. Compared to the Single Localization No Spoofing (SL-NS) case, where vehicles relied solely on their own sensors for localization, the CL-NS case shows:

- *Improved Error Distribution:* The points in the Stanford diagrams for all vehicles are more tightly clustered near the origin, indicating a significant reduction in localization error (see Figure 9). This grouping reflects that most localization errors are now well below 2 meters, with no data points exceeding the threshold that would place them in the MI or HMI zones.
- *Decreased Mean Error:* The average error for each vehicle has decreased due to the cooperative approach, meaning that the shared position data allows each vehicle to correct its position estimate more effectively, leading to better precision for the entire fleet.

The absence of points in the MI and HMI zones across all vehicles in the cooperative case is a direct indication of the system's improved integrity with respect to the single localization case. By sharing data, the vehicles are better equipped to avoid undetected large errors.

The system's ability to avoid MI and HMI conditions through cooperation highlights the critical role of shared localization in enhancing the integrity and reliability of autonomous navigation systems, especially in complex environments where single localization may fail or degrade.

4. Cooperative Localization Spoofing (CL-S)

In this case, we analyze the performance of Cooperative Localization while Car 1 is spoofed. To further explore the impact of spoofing, we introduce variations in the measurement noise on the relative observation, analyzing how the spoofing affects performance when a vehicle either has higher confidence in its own solution or relies more heavily on relative observations.

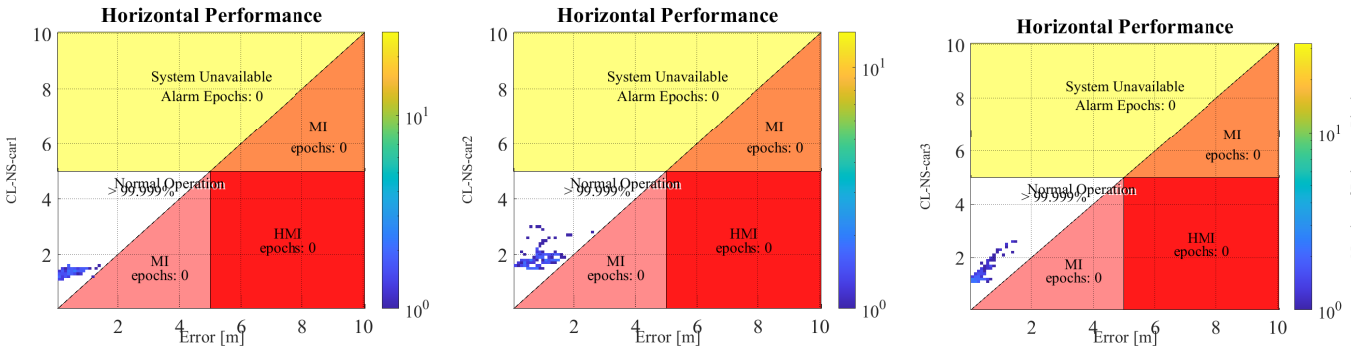

(a) Car 1 **(b)** Car 2 **(c)** Car 3

Figure 9: Integrity monitoring for the case CL-NS

a) Higher Confidence in Vehicle's Own Solution

In this scenario, we analyze the system performance when the relative observation noise is high, meaning that the vehicle relies more on its own positioning solution and less on the relative observations from other vehicles.

In this scenario, Car 1 still performs poorly, with 138 epochs in the HMI zone, similar to the SL-S case (see Fig. 11). Normal operation remains comparable, but there are 13 epochs in the System Unavailable zone, slightly better than the 16 epochs in SL-S. While the overall performance of Car 1 remains poor, this indicates that cooperative localization has a modest positive impact, even when Car 1 places more reliance on its own (compromised) positioning solution. This small improvement is due to the partial benefits of cooperation, which help mitigate some, though not all, of the spoofing impact.

Car 2, in this case, has 59.76% of points in the Normal Operation region, with 103 epochs in the MI zone. This is worse than in the SL-S case, where 66.77% of points were in Normal Operation. This degradation in performance is expected, as Car 2, when relying more on its own solution rather than on cooperative data, becomes more susceptible to indirect influences from the spoofed vehicle (Car 1), though it is still not as severely impacted as in a scenario with more trust in relative observations.

Car 3 continues to perform exceptionally well, with 99.21% of points in Normal Operation, only 2 epochs in the MI zone, and no points in the HMI or System Unavailable zones. This is similar to its performance in SL-S. The minor drop in performance compared to CL-NS (with almost no MI epochs in CL-NS) suggests that Car 3's standalone localization is strong enough to handle the situation, with cooperative localization adding minimal degradation even with spoofing.

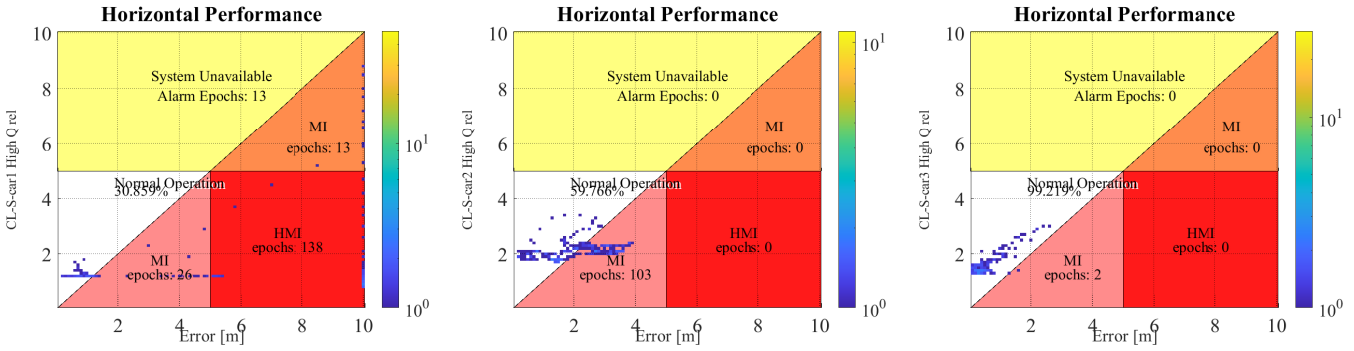

(a) Car 1 **(b)** Car 2 **(c)** Car 3

Figure 10: Integrity monitoring for the case CL-S with lower confidence in relative observation

b) Higher Confidence in Relative Observation

In contrast, we explore the case where relative measurement noise is low, indicating that the vehicle places more trust in the relative observations from other vehicles, effectively relying on the cooperative localization system.

In this scenario, Car 1 shows a slightly better performance than when it relied more on its own position solution. With 136 epochs in the HMI zone and 74 epochs in the MI zone, this is an improvement compared to the high relative observation noise scenario (where it had 138 HMI epochs). Importantly, the System Unavailable epochs have decreased. This shows that by placing more

trust in relative observations, cooperative localization helps Car 1 avoid system unavailability and slightly improves its overall performance, even though it continues to struggle with localization errors due to spoofing.

For Car 2, performance degrades considerably. With only 36.71% of points in Normal Operation and 162 epochs in the MI zone, this is much worse than both the High measurement noise and SL-S scenarios. This indicates that when Car 2 places more trust in cooperative data, it becomes more susceptible to the spoofed data from Car 1, leading to severe localization issues. This is an expected outcome given that trusting spoofed vehicles in a cooperative system propagates errors more easily.

Car 3 remains largely unaffected, with 99.21% of points in Normal Operation and only 2 epochs in the MI zone, similar to its performance in both SL-S and High relative measurement noise scenarios.

The key takeaway is that Car 3's localization performance is robust, and even under spoofing, cooperation helps it avoid critical integrity failures. The reason behind this behaviour is still under investigation but it would seem that vehicle 3 is more robust because of the choice of trajectory with an environment more conducive to optimal reception of GNSS observations. A more in-depth study of the elevations of the satellites received and the CN0 ratios and other metrics will reveal a little more about this phenomenon of non-homogeneous resilience.

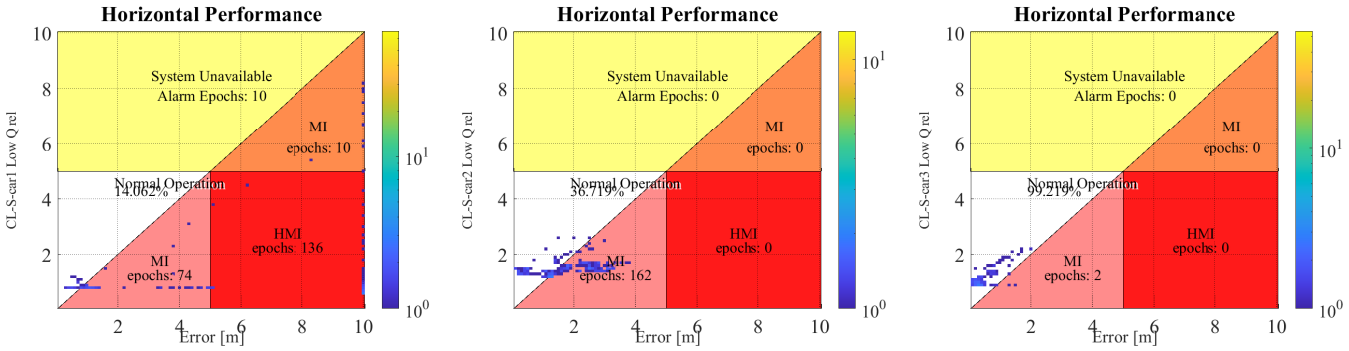

(a) Car 1 **(b)** Car 2 **(c)** Car 3

Figure 11: Integrity monitoring for the case CL-S with higher confidence in relative observation

5. Result discussion

The results from the various cases highlight the significant influence of both spoofing and cooperative localization on the performance of the fleet. When spoofing is introduced in the SL-S case, the performance of Car 1 deteriorates drastically, with its points extending into the HMI zone. Despite the rest of the fleet (Cars 2 and 3) remaining unaffected in this single localization scenario, the lack of data-sharing makes Car 1 highly vulnerable to spoofing.

In the CL-NS case, we see a marked improvement in performance for all vehicles, especially Car 2, whose weaker performance in the single localization scenario is significantly improved. Cooperation brings the vehicles' errors closer to the origin, minimizing overall localization errors. However, in the CL-S case, while Car 1 still suffers from the spoofing attack, its performance improves compared to SL-S due to the support of cooperative localization. Nonetheless, Car 2, which performed well in the CL-NS case, now shows degradation with some points crossing into the MI zone, and a few points entering the HMI zone. This indicates that cooperative systems, while improving robustness, can also propagate erroneous data when one vehicle is compromised, particularly for vehicles with weaker initial performance in the single localization mode.

Table1 summarizes the Root Mean Square Error (RMSE) for the X and Y errors of each vehicle across the different cases. In

Case	X Error (Car 1)	X Error (Car 2)	X Error (Car 3)	Y Error (Car 1)	Y Error (Car 2)	Y Error (Car 3)
SL NS	1.3202	2.4655	0.4945	0.7527	2.2272	1.2024
SL S	12.3947	2.4655	0.4945	3.5313	2.2272	1.2024
CL NS	0.5983	1.1702	0.3184	0.2999	1.0256	0.5135
CL S Q_{rel}^{High}	12.4154	2.3100	0.5221	3.3908	2.1471	1.2144
CL S Q_{rel}^{Low}	12.4300	2.1621	0.6010	3.2362	2.0940	1.1944

Table 1: Comparison of X and Y Mean Square Root Errors for Car 1, Car 2, and Car 3

the SL-NS case, Car 1 shows moderate error, with an RMSE of 1.3202 meters in the X direction and 0.7527 meters in the Y

direction, reflecting decent performance. Under SL-S, where Car 1 is spoofed, the RMSE jumps to 12.3947 meters (X) and 3.5313 meters (Y), indicating a severe localization error due to the spoofing attack. In the CL-NS case, Car 1 benefits from cooperation, with RMSE values dropping to 0.5983 meters (X) and 0.2999 meters (Y), demonstrating how cooperation across the fleet improves positioning accuracy.

In the CL-S case, where Car 1 is spoofed but can rely on data from other vehicles, the RMSE remains high. With high Q_{rel}, Car 1's RMSE is 12.4154 meters (X) and 3.3908 meters (Y), showing a clear degradation compared to the CL-NS case. With low Q_{rel}, where Car 1 places more trust in the relative observations, its RMSE is slightly better at 12.4300 meters (X) and 3.2362 meters (Y). Although cooperation mitigates some of the spoofing's impact, the attack still degrades Car 1's performance, particularly in the Y direction compared to the non-spoofed cooperative case.

In the SL-NS case, Car 2 experiences higher errors compared to Car 1 and Car 3, with an RMSE of 2.4655 meters (X) and 2.2272 meters (Y). This indicates that Car 2's performance is relatively weaker in single localization without spoofing, and it keeps the same errors in the SL-S case, as it has not been spoofed or interacted with the spoofed vehicle (Car 1). In the CL-NS case, cooperative localization improves Car 2's performance, with the RMSE dropping to 1.1702 meters (X) and 1.0256 meters (Y), showing that cooperation helps correct the errors for weaker vehicles by leveraging data from the fleet.

However, in the CL-S case, Car 2's RMSE increases when exposed to spoofed data from Car 1. With high Q_{rel}, Car 2's RMSE increases to 2.3100 meters (X) and 2.1471 meters (Y), while with low Q_{rel}, it slightly improves to 2.1621 meters (X) and 2.0940 meters (Y). While the X error does not change, the Y error suggests that the shared spoofed data from Car 1 still negatively impacts Car 2's localization, highlighting the vulnerability of cooperative systems to data corruption.

Car 3 consistently performs the best in all cases, showing relatively low errors. In the SL-NS case, the RMSE is 0.4945 meters (X) and 1.2024 meters (Y), indicating strong performance in single localization without spoofing. Similarly, in the SL-S case, Car 3's performance remains unchanged and unaffected by the spoofing of Car 1. In the CL-NS case, cooperative localization further improves Car 3's performance, reducing the RMSE to 0.3184 meters (X) and 0.5135 meters (Y), demonstrating the benefits of cooperation even for well-performing vehicles.

In the CL-S case, Car 3's RMSE increases slightly due to the spoofed data. With high Q_{rel}, the RMSE increases to 0.5221 meters (X) and 1.2144 meters (Y), and with low Q_{rel}, relying therefore more on relative observation, it increases to 0.6010 meters (X) and 1.1944 meters (Y). The X error increases slightly more in the low Q_{rel} case, reflecting a degradation in performance due to the impact of the spoofed vehicle's shared information. However, Car 3's overall performance remains strong, and the increase in error is minimal compared to Car 2. This shows that Car 3's robust localization system is less affected by the spoofed data and continues to perform well even when relying more on cooperative data.

V. CONCLUSION AND PERSPECTIVES

This study investigated the performance of cooperative localization in a vehicle fleet, particularly focusing on the effects of spoofing on one vehicle and its potential to impact the entire network. The findings demonstrate that cooperative localization significantly enhances the overall accuracy of the fleet's positioning by reducing localization errors, especially for vehicles that initially exhibit weaker performance in single localization scenarios. However, the results also reveal that cooperation introduces certain vulnerabilities. In the case of a spoofing attack on a single vehicle, the compromised data can propagate through the fleet and negatively affect other vehicles, as observed in the slight degradation of Car 2's performance.

Nevertheless, the spoofed vehicle benefited from cooperative localization. In comparison to the single localization spoofing scenario, the spoofed vehicle experienced a reduction in system unavailability and hazardous misleading information. This highlights the potential of cooperation to dilute the impact of spoofing and maintain system integrity under adverse conditions, although it is not immune to errors being propagated within the network. The overall improvement in performance for non-spoofed vehicles, particularly Car 3, underlines the robustness of cooperative systems, though the negative effect observed in Car 2 points to a need for enhanced detection and mitigation mechanisms to prevent error propagation from compromised vehicles.

Moreover, it can be argued that cooperation not only enables a form of shared detectability but also provides a lever for mitigation that is absent in standalone positioning. If one or even several vehicles are spoofed, either they can detect it on their own, or the fleet as a whole can assist in detecting the compromised vehicles. This allows for dynamic adjustments in the confidence attributed to each observation. Spoofed vehicles should place greater confidence in the cooperative information from non-spoofed vehicles while reducing their reliance on both their standalone data and the cooperative data from other potentially spoofed vehicles. Conversely, non-spoofed vehicles will maintain high confidence in their own standalone positioning and in data from vehicles identified as non-spoofed, enabling more effective mitigation strategies.

The limitations of this study arise primarily from the reliance on simulations based on Ultra-WideBand (UWB) technology for vehicle distance estimation, which may not fully capture the complexities of real-world conditions. In future research, we plan to test with real sensor-based acquisition solutions that support cooperation to assess the practical implications in a dynamic

environment.

One of the most important factors in mitigating spoofing attacks is the speed at which the system can detect compromised vehicles and alert the fleet. In the current study, we focused on evaluating the impact of spoofing without implementing detection mechanisms. Future studies should explore advanced detection algorithms that can minimize the delay between the onset of an attack and the system's response.

Another important direction for future research is the exploration of tightly coupled cooperative localization systems. The current study used a loosely coupled approach, where vehicles shared processed data. A tightly coupled system, in which raw GNSS data is shared between vehicles, could offer greater precision and improved robustness against spoofing attacks. In such systems, vehicles would have access to more granular information, enabling them to more accurately assess and correct positioning errors.

Further mitigation strategies involve integrating additional information sources to cross-check the accuracy of position estimates. Using multiple prediction models could also help detect when a vehicle's position deviates from expected trajectories, identifying inconsistencies that may suggest a spoofing attack in various conditions.

Finally, information-theory-based approaches present an avenue to explore for improving spoofing detection. These methods could help create fault-tolerant solutions and identify fault indicators, which would assist in localizing faulty sensors and satellites, thereby strengthening the overall integrity of the cooperative localization system.

ACKNOWLEDGEMENT

This work benefits from the financial support of the ANR French National Research Agency within the project LOCSP No 2019-CE22-0011. The authors thank PRETIL platform for the disposal of the equipment used for this study.

REFERENCES

D. Akos. Who's afraid of the spoofer? gps/gnss spoofing detection via automatic gain control (agc). *Annual of Navigation*, 59: 281–290, 2012. doi: 10.1002/NAVI.19.

P. A. Boysen and H. Zunker. Integrity hits the road. *GPS World*, 16(7):30–, 2005.

A. Broumandan and G. Lachapelle. Spoofing detection using gnss/ins/odometer coupling for vehicular navigation. *Sensors*, 18 (5):1305, 2018. doi: 10.3390/s18051305.

S. Dasgupta, M. Rahman, M. Islam, and M. Chowdhury. A sensor fusion-based gnss spoofing attack detection framework for autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 23(12):23559–23572, Dec 2022. doi: 10.1109/TITS.2022.3197817.

G. Dherbomez, Z. El Mawas, M. Duquesne, G. De Smet, C. Cappelle, and M. El Badaoui El Najjar. Multi-vehicles dataset for ion gnss+ 2024 conference. Recherche Data Gouv France, 2024. doi: 10.57745/3C63J0.

I. Fernendez-Hernendez et al. Increasing international civil aviation resilience: A proposal for nomenclature, categorization and treatment of new interference threats. In *2019 International Technical Meeting of The Institute of Navigation*, pages 389–407, Reston, Virginia, 2019. doi: 10.33012/2019.16699.

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

C. Kang, S. Y. Kim, and C. G. Park. Adaptive complex-ekf-based doa estimation for gps spoofing detection. *IET Signal Processing*, 12:174–181, 2018. doi: 10.1049/iet-spr.2016.0646.

Z. El Mawas, C. Cappelle, M. Daher, and M. E. B. El Najjar. Comparative analysis of centralized and federated learning techniques for sensor diagnosis applied to cooperative localization for multi-robot systems. *Sensors*, 23(17):7351, 2023. doi: 10.3390/s23177351.

J. R. V. D. Merwe, X. Zubizarreta, I. Lukcin, A. Rugamer, and W. Felber. Classification of spoofing attack types. In *2018 European Navigation Conference (ENC)*, pages 91–99, Gothenburg, Sweden, 2018. IEEE. doi: 10.1109/EURONAV.2018.8433227.

C. Pierre, R. Chapuis, R. Aufrère, J. Laneurit, and C. Debain. Range-only based cooperative localization for mobile robots. In *2018 21st International Conference on Information Fusion (FUSION)*, pages 1933–1939, 2018. doi: 10.23919/ICIF.2018.8455692.

M. L. Psiaki, B. W. O'Hanlon, J. A. Bhatti, D. P. Shepard, and T. E. Humphreys. Gps spoofing detection via dual-receiver correlation of military signals. *IEEE Transactions on Aerospace and Electronic Systems*, 49(4):2250–2267, 2013. doi: 10.1109/TAES.2013.6621814.

F. Rothmaier. *Statistical Inference for Safe and Continuous Navigation in the Presence of GNSS Spoofing*. PhD thesis, Stanford University, USA, 2021.

E. Shafiee, M. Mosavi, and M. Moazedi. Detection of spoofing attack using machine learning based on multi-layer neural network in single-frequency gps receivers. *Journal of Navigation*, 71:169–188, 2017. doi: 10.1017/S0373463317000558.

H. Shahid, L. Canzian, C. Sarto, O. Pozzobon, J. Reyes-Gonzalez, G. Seco-Granados, and J. López-Salcedo. Statistical characterization of snapshot osnma spoofing detection. In *2023 International Conference on Localization and GNSS (ICL-GNSS)*, pages 1–6, 2023. doi: 10.1109/ICL-GNSS57829.2023.10148915.