GNSS-WASP: GNSS Wide Area SPoofing

Christopher Tibaldo, Harshad Sathaye, Giovanni Camurati, Srdjan Capkun
Department of Computer Science
ETH Zurich, Switzerland

Abstract

In this paper, we propose GNSS-WASP, a novel wide-area spoofing attack carried by a constellation of strategically-located synchronized transmitters. Unlike known attacks, which are constrained by the attacker's ability to track victim receivers, GNSS-WASP manipulates the positions measured by all the receivers in a target area without knowing the victim's positions. This allows GNSS-WASP to spoof a swarm of victims to another location while preserving their true formation (i.e., their relative distances). This opens the possibility of advanced attacks that divert entire fleets of vehicles and drones in a large area without the need to track specific victims. As such, GNSS-WASP bypasses state-of-the-art spoofing countermeasures that rely on constellations of receivers with known distances and those that rely on sudden, unpredictable movements for spoofing detection. While previous works discuss the stringent requirements for perfect spoofing of multiple receivers at known fixed locations, GNSS-WASP demonstrates how to spoof any number of moving receivers at unknown positions in a large area with an error that can remain hidden behind the legitimate noise. In addition to extensive simulations, we implement a prototype of GNSS-WASP with off-the-shelf software-defined radios and evaluate it on real GNSS receivers. Despite the error introduced by the proposed attack, GNSS-WASP can successfully spoof two receivers while maintaining their relative distance with an average error of 0.97 m for locations 1000 m away from the reference position. Finally, we also highlight possible countermeasures.

1 Introduction

Global Navigation Satellite Systems (GNSS) provide worldwide geo-localization and time-synchronization services in real time. Multiple constellations and services such as GPS, Galileo, GLONASS, and BeiDou have been developed by different countries, providing excellent coverage and accuracy.

A major threat to the security of applications using GNSS is the feasibility of spoofing attacks, capable of altering the position of a victim receiver.

In its simplest form, a spoofing attack consists of a transmitter broadcasting the combined satellite signals visible at a desired location, making any receiver in range of the attack believe it is situated at this location [32]. In recent years, reports of these attacks in the wild have increased, such as in the world of aviation, where aircraft pilots have reported an increasing number of suspected GPS spoofing incidents in various areas of the world [27, 45] posing a real threat to the safety and security of travelers.

Multiple countermeasures aiming at detecting spoofing attacks have been proposed [50, 55] but no fundamental solution to the problem exists. Since the position is calculated from the arrival time of the signals broadcast by the satellites, cryptographic authentication of the navigation messages is generally insufficient to protect against spoofing. For example, Montallebighomi et al. [41] have recently shown how to bypass Galileo Open Service Navigation Message Authenication (OSNMA) by selectively delaying satellite signals. However, some scenarios exist in which spoofing attacks are particularly challenging and prone to detection.

The first such scenario involves spoofing a group of multiple receivers simultaneously while preserving the relative distances of a set of receivers. Tippenhauer et al. [62] show that this *group spoofing* problem severely constrains the attacker in terms of number of transmitters required and/or places at which they should be located. The complexity of the attack exponentially increases as the number of receivers increases. To further increase the complexity, such an attack is specific to a set of fixed victims. Hence, checking the consistency of the relative positions between multiple receivers has been proposed as an effective countermeasure [34, 59].

The second challenging scenario involves spoofing a moving receiver. The attacker has to track the victim and constantly update the spoofing signal. Sathaye et al. [52] propose to detect spoofing by checking the consistency between GNSS and inertial sensors while performing a sudden, unpredictable movement. Tippenhauer et al. [62] observe that movements are also a challenge for the group spoofing problem, as they change the constraints on the attacker positions.

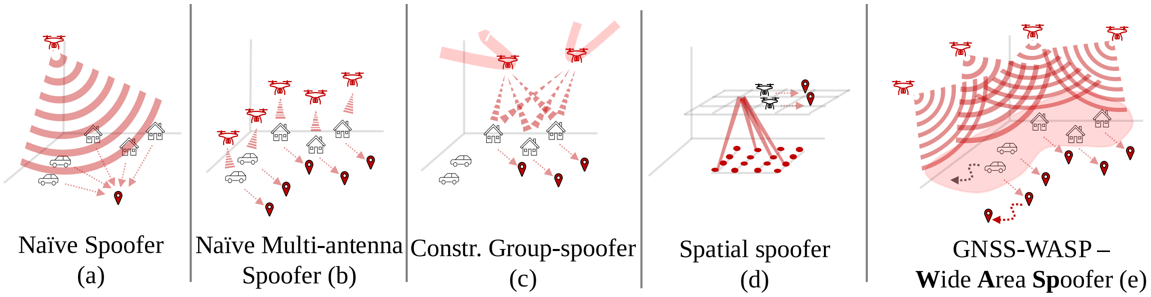

Naïve Spoofer (a)

Naïve Multi-antenna Spoofer (b)

Constr. Group-spoofer (c)

Spatial spoofer (d)

GNSS-WASP – **W**ide **A**rea **Sp**oofer (e)

Figure 1: **`GNSS-WASP` vs. existing attack strategies** - see Table 1 for details.

		Single antenna (a) [24, 32, 53, 62]	**Multi antenna (b)** [62]	**Group spoofer (c)** [62]	**Spatial (d)** [20]	**`GNSS-WASP` (e)** *This paper*
Requirements	**Target**	Victim(s) in range	Individual victims	Individual victims	2D area above	3D area
	Transmitters	1	One per victim	One per sat. (>= 4)	Array size	One per sat. (>= 4)
	Antenna type	Omnidirectional	Directional	Directional	Antenna array	Omnidirectional
	Synchronization	n.a.	Good	Tight	Tight	Tight
	Attacker Placement	Any position in range	Close to each individual victim	Constrained by victims positions	Ground below attack surface	In line with satellites
Capabilities	**Spoofed positions**	Same for all	One per victim	One per victim	One per victim	One per victim
	Position updates	By attacker	By attacker	By attacker	By attacker/victim	By attacker/victim
	Relative distances	Violated	Preserved for target	Preserved for target	≈ Preserved in area	≈ Preserved in area
	Movements	Preserved w/ update	Preserved w/ update	Preserved w/ update	≈ Preserved in area	≈ Preserved in area
Countermeasures	**Multiple Rx [34]**	High detection	Low detection	Low detection	Low detection	Low detection
	Sensor fusion [52]	Bypassed w/ position update	Bypassed w/ position update	Bypassed w/ position update	Low detection	Low detection
	Exit area	n.a.	n.a.	n.a.	Bypassed w/ update	Bypassed w/ update
	Approximation error in area	n.a.	n.a.	n.a.	Low detection High latency	Low detection High latency
Summary	**Main strength**	Simple design	Flexible	Scalable	2D area wide	3D area, scalable
	Main limitation	Inconsistent positions	One antenna per victim	Constrained by victim positions	Wrong angle, 2D Antenna array	Approximation error
	Implementation	Many E.g., [24, 32, 53]	Would use multiple (a)	Math model at fixed position	Mathematical model only	Hardware prototype Numerical model

Table 1: **Attack strategies**. Two defenses significantly increase spoofing complexity: (i) checking the consistency of the positions of multiple receivers placed at known distances [34], and (ii) verifying the consistency of position and inertial sensors during (unpredictable) movements [52]. The first constrains the number (b) or position (c) of the attack transmitters [62]. The second forces the attacker to update its signal and (constrained) locations based on fine-grained tracking of the victims. To relax these constraints, (d) [20] approximates the legitimate signal on a 2D surface using an antenna array on the ground and a Wiener filter, but it has limited coverage and transmits from an easily-detectable angle. An ideal attacker would use the same locations as the legitimate satellites. `GNSS-WASP` (e) places its transmitters in line with the satellites at lower altitudes. By construction, this creates a 3D area where the ideal attack is approximated with a small error, and any victim is spoofed consistently. The higher the attackers, the larger the 3D area. Outside the area, (e) converges to (a). A victim can detect (d, e) by exiting the area unless the attacker updates its coverage. Distinguishing the approximation error of (d, e) from noise requires many measurements and high latency. Of the advanced strategies, only (e) has a prototype and was tested on both software-defined and commercial receivers

If these countermeasures are in place, simply impersonating the entire constellation with a single transmitter would not be sufficient. Successful spoofing would require either attacking each victim individually with directional transmitters, which is not practical, or placing the attacker in a very constrained set of locations [62], which would have to be re-calculated and updated following the victim's movements. To relax these constraints and make the spoofing of drone swarms more practical, Ceccato et al. [20] propose to approximate the spoofing signal for all locations on a 2D surface, using an antenna array placed on the ground and a Wiener filter. However, coverage is limited to a plane and requires many antennas to reduce the error. Moreover, the incorrect angle (signal originating from the ground) is detectable. None of these advanced strategies [20, 62] have been demonstrated beyond a mathematical model.

In this paper, we present GNSS-WASP, a novel GNSS spoofing attack that consistently manipulates the position of any group of receivers in a target area. The consistency of relative positions and movements of the receivers is preserved, albeit with a minimal error hidden in natural noise. GNSS-WASP can thus bypass state-of-the-art countermeasures [34, 52]. A small fixed number of attack devices (at least four) strategically placed above the target area can spoof an arbitrary number of victims. The attack assumes that the victims are within the target area and does not require prior knowledge of the precise location of the victims or adjustments based on tracking their movements within this space. Depending on the attack scenario, an attacker can still move the effective attack area to ensure victims stay inside its bounds, making GNSS-WASP even more effective.

Successful group spoofing can be achieved by relaxing the idealized constraints set by Tippenhauer et al. [62]. Instead of analyzing how perfect GNSS receivers limit the positioning of attacking transmitters we identify a transmitter geometry which minimizes errors over a wide area. Figure 1 and Table 1 compare GNSS-WASP with existing attack strategies and requirements [20, 62].

We investigate possible GNSS-WASP-aware countermeasures and show that, although this attack successfully thwarts detection strategies based on relative distances or movement of victims, GNSS-WASP can be detected by a dedicated receiver array that continuously collects error statistics. The trade-off for these countermeasures is their increased complexity and time requirement. Thus higher latency and cost is incurred in detecting and validating the measured position.

We provide a theoretical model that explains the working principle of GNSS-WASP and analyze the errors within a given target area. We extensively simulate different attack scenarios on the open-source GNSS-SDR receiver and evaluate them against state-of-the-art countermeasures. Finally, we implement a prototype of GNSS-WASP with off-the-shelf software-defined radios and perform tests using commercial GPS receivers. We show that GNSS-WASP can achieve an average accuracy of up to 10.52 m in the derived position of a single receiver against the ground-truth for locations up to 750 m from the reference position. Even though the error to the ground-truth increases as the receivers move further away from the reference position, the relative distance between multiple nearby receivers stays close to the ground truth. GNSS-WASP can spoof two receivers located at approximately 10 m distance from each other, while maintaining their relative distance with an average error of 0.97 m for locations 1000 m away from the reference position. In short, we make the following contributions:

- We propose GNSS-WASP, a novel wide-area GNSS spoofing attack, capable of consistently manipulating the position of an arbitrary number of victim receivers in a target area using a limited number of transmitters, without knowing the victims' precise locations and movements.
- We develop GNSS-WASP with a novel approach to group spoofing. Instead of studying the requirements for spoofing a specific group of receivers at known locations, we analyze how a limited number of strategically placed attack transmitters affect the error observed by receivers within a targeted area.
- We implement GNSS-WASP in simulation and with a prototype based on off-the-shelf software-defined radios synchronized with GPS Disciplined Oscillators (GPSDOs). We extensively evaluate GNSS-WASP in multiple scenarios and against existing countermeasures. For instance, GNSS-WASP can successfully spoof two receivers while maintaining their relative distance with an average error of 0.97 m for locations 1000 m away from the reference position. Additionally, through extensive simulations, we show that GNSS-WASP can successfully spoof a group of four receivers arranged in a square formation in a 1 km^2 area without being detected.
- We discuss possible mitigation strategies that can effectively detect the attack based on the error introduced by GNSS-WASP. Strategies such as measurements over a 1 s rolling window and cumulative sum (CUSUM), along with widely spaced receiver arrays, significantly improve detection rates.

2 Background

Satellite Navigation Overview

GNSS signals use the Direct Sequence Spread Spectrum (DSSS) technique as a multiple access scheme. Each satellite vehicle is assigned a number to generate a pseudorandom noise (PRN) sequence of 1s and 0s. The satellite uses this sequence to spread and transmit data that enables the receiver on the ground to i) obtain the current time and ii) calculate the satellite's position at the time of transmission of a particular message. It is important to note that the algorithm to generate the PRN sequence has been made publicly available. These messages are transmitted on frequencies within the L-band (1 GHz to 2 GHz); for example, GPS uses L1, L2, and L5 frequencies. In this work, we primarily focus on civilian GPS signal transmitted on L1 frequency.[1] Upon start-up, a receiver performs a 2D search of a particular PRN sequence in the time and frequency domain. Next, the receiver tracks the carrier using the time delay and Doppler estimated in the earlier step and calculates the pseudo-range between itself and the satellites. Finally, the receiver calculates its own location once it successfully estimates ranges to at least four satellites.

GNSS Attacks Overview

Satellite navigation systems are open systems by design; their modulation schemes, spreading sequences, and signal processing are public knowledge. This enables attackers to generate fake signals to disrupt legitimate GNSS signals. Broadly, there are two types of attacks on GNSS: 1) Jamming attacks, where an adversary transmits high-power noise to jam or block legitimate signals and deny GNSS services, and 2) Spoofing attacks, where an adversary transmits pre-crafted signals and forces all the receivers in its vicinity to calculate a location of its choosing. Since GNSS signal construction is public knowledge, there are several commercial systems that are capable of generating signals from different satellite navigation systems [3, 7, 9]. Moreover, open-source software like GPS-SDR-SIM [24] and Galileo-SDR-SIM [53] and inexpensive Software Defined Radio (SDR) [1, 4, 26] allow users to generate and transmit fake GNSS signals to deceive receivers at a very low cost.

The receiver uses the *common reception time* or the *common transmission time* techniques [39] for pseudorange calculation; a receiver can be deceived even by attackers with a single antenna. Thus, most of the attacks described in the literature so far consider an attacker with a single antenna. In this attack type, the attacker generates signals from all the visible satellites, adjusts their relative offsets, and transmits a combined signal. However, such attacks can be easily detected by leveraging spatially diverse observations of the received signal. Receivers capable of motion can also use sensor fusion techniques [22, 38, 61] to detect spoofing attacks, as an attacker is required to constantly track the victim receiver to generate an appropriate spoofing signal. Works by Tippenhauer et al. [62] and Jansen et al. [34] introduce a countermeasure where they use multiple co-located receivers to detect a single antenna attacker. In theory [62] introduces an attacker model that is capable of spoofing multiple co-located receivers. However, such an attacker is heavily dependent on the location of the victims, and the complexity further increases if the victims move. An improved attack that approximates the spoofing signal for all locations on a 2D surface was proposed in [20], but it requires an antenna array on the ground. Refer to Figure 1 and Table 1 for a detailed comparison of different attacker strategies and GNSS-WASP. Furthermore jamming can be employed in combination with spoofing, to avoid sudden jumps in receiver position as well as to prevent re-acquisition of authentic signals during an attack. The weak authentic signal is hereby overshadowed by the combination of the attacker's stronger signal and noise [32].

3 GNSS-WASP

GNSS-WASP is a novel spoofing attack strategy that manipulates the position of any number of receivers in a target area, preserving their relative distances and movements at the net of a small error hidden in the noise. In this section, we highlight the threat model, explain the attack strategy, and provide a numerical model describing the intuition of our attack strategy.

3.1 Threat model

Contrary to prior works focusing on an attacker equipped with a single transmitter, GNSS-WASP uses strategically placed synchronized transmitters that transmit a generic *non-personalized* spoofing signal. Unlike single-antenna attackers, GNSS-WASP does not rely on prefabricated temporal alignment of the attack signal; instead, it leverages the spatial diversity of its victims to set the appropriate temporal alignment of the arriving signal, just like a legitimate satellite navigation system does. This requires the attacker to be able to place at least four transmitters in specific locations. For this task, we propose using drones equipped with RTK GPS for accurate, centimeter-level positioning [8, 23, 28]. We assume that the attack starts with a cold start or a smooth takeover, and that the attack signal overshadows the legitimate transmissions, burying the signals from all legitimate satellites in noise. Moreover, just like the legitimate GNSS constellation, our attacker must also be able to synchronize the clocks on its transmitters at the nanosecond level, which is possible using GPS Disciplined Oscillators [10].

Since GNSS-WASP leverages the channel for temporal alignment of the spoofing signal, an attacker can affect a wide area without knowing the precise location of its victims. However, it still needs to know the approximate location of its target. This is vital in ensuring proper coverage of the target area. Next, we assume that the attacker has access to legitimate satellites' orbital information, as GNSS-WASP needs to closely follow the trajectory of the satellite that it is impersonating. We do not require the attacker to precisely track satellite trajectories since satellite trajectories are known and ephemeris data are publicly available [44]. In addition, since attackers impersonate the satellites from a much lower altitude close to the Earth, their position has to be updated at a very low speed that drones can easily achieve.

[1] L1 = 1.57542 GHz

3.2 Attack Scenarios

The proposed attack is most effective when the attacker needs to spoof multiple receivers simultaneously without personalized spoofing signals and spoof receivers that are in motion without having to track them precisely. Following are some use-cases where GNSS-WASP is effective:

Co-located Receivers: In [62] authors give an example of a ship with multiple receivers placed such that the distance between each other is greater than the inherent noise of GPS receivers. In the case of a single antenna attacker, to successfully spoof such a vessel, the attacker will have to individually spoof each receiver. This is necessary to ensure the receivers maintain their relative positions and preserve the constellation's geometry, even after spoofing. Here, GNSS-WASP can effectively spoof without precisely tracking the vessel and without individually spoofing each receiver.

Swarm of Autonomous Vehicles: Works like [31, 47, 54, 56] have shown the ability to hijack a single UAV by attacking their GPS receivers. However, these works focus on attacking a *single* UAV. The popularity of UAV swarms is rising, especially in surveying and search and rescue operations [12]. Moreover, works like [42] use UAV swarms to detect spoofing collectively. Since GNSS-WASP leverages the channel and spatial diversity, the swarm will maintain its formation, allowing an attacker to gain complete control of the swarm without knowing the precise location of each UAV. In this case, an attacker will need to know the general area of the target swarm and place its transmission antennas at a suitable distance.

INS/GNSS Spoofing Detection: Inertial Navigation Systemss (INSs) based on inertial sensors such as accelerometers and gyroscopes play a vital role in maintaining stability in modern vehicles. However, they are highly accurate but perform poorly to provide long-term stability. Thus, they are fused with GNSS measurements to get a more robust position estimate. Strategies outlined in [22, 38, 52, 61] compare inertial measurements and GNSS measurements to detect GNSS spoofing attacks. Such techniques are very effective in detecting attacks. When a receiver receives signals from GNSS-WASP, the computed location is governed by the receiver's motion and there-by with the changing channel between the transmitter and the receiver. This ensures that the computed locations automatically align with the inertial measurements, thus defeating any security checks as long as the victims remain within the target area. Should the victim head outside the area, the attacker could coarsely update the position of the target area to ensure they stay inside.

Hidden Receivers: It is a common practice to install multiple continuously operating ground reference stations to monitor the integrity of GNSS signals within large facilities like airports or industrial complexes. These receivers can detect naive spoofing attacks provided there is sufficient spatial separation between these receivers [34]. To execute an undetected spoofing attack, an attacker will have to spoof each reference station individually, requiring the attacker to know the exact location of each receiver. In such a scenario, GNSS-WASP can spoof these receivers without knowing their individual locations while still maintaining their relative distances.

3.3 Attack Overview

GNSS-WASP leverages a multi-transmitter setup to spoof receivers distributed in a wide area. Refer to Figure 2 for a graphical overview of the attack. Unlike previous spoofing approaches that involve fabricating temporal alignment of satellite signals, GNSS-WASP uses the channel to achieve temporal diversity as legitimate satellites. GNSS-WASP uses a constellation of flying drones equipped with transmitters. These transmitters are synchronized to ensure phase and frequency alignment. Each drone 'shadows' a satellite, i.e., it closely follows the satellite's orbit it is spoofing, thus replicating an actual GNSS satellite constellation. The drones are positioned along the axis that runs through the ideal satellite position and a reference location R as depicted in Figure 2. Next, it generates and transmits a signal such that the pseudorange calculated at the drone's location is y (See Figure 3). Since signals from all attacker transmitters are synchronized, all receivers located within the radio range of the attacker transmitters can calculate a unique PVT solution depending on their location relative to the attacker, just as they would with legitimate GNSS satellites.

Ideally, the attacker should launch at least four transmitters since that is the minimum number of satellites required to estimate the position. An attacker can create a more realistic illusion by increasing the number of transmitters and thus the number of shadowed satellites. An attacker may even spoof satellites from multiple constellations. The main limiting factor is the ability to deploy and coordinate many drones. When receivers are located at the attack's reference location, the signal generated by the attackers will be indistinguishable from an actual signal. However, if the receiver is not present on the axis of the attacker transmitter and the satellite, it experiences a small systematic error, explained further in Section 3.4. As receivers move further from the center of the attack, the magnitude of these systematic errors increases in the observed position. By moving transmitters further away from the reference position, an attacker can expand the area where the magnitude of the systematic error stays small, making detecting the attack more difficult.

3.4 Numerical Model

We provide a simple mathematical model to explain the working principle of GNSS-WASP and study the approximation errors it introduces.

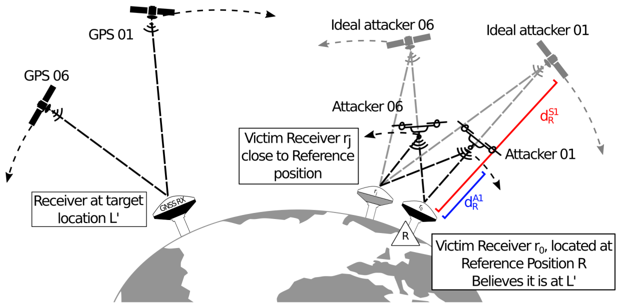
GPS 01
GPS 06
Ideal attacker 06
Ideal attacker 01
Attacker 06
Attacker 01
Victim Receiver r_j close to Reference position
Receiver at target location L'
GNSS RX
r_j
r_0
R
d_R^{S1}
d_R^{A1}
Victim Receiver r_0, located at Reference Position R Believes it is at L'

Figure 2: **Overview of `GNSS-WASP`.** The goal is to convince all receivers near the *reference position R* that they are instead near the *target location L* while preserving the consistency of their relative positions and movements. An *ideal attacker* would replicate over R the legitimate constellation visible over L. `GNSS-WASP` replicates the legitimate constellation approximately by placing the transmitters A_i at a constant distance $d_R^{A_i}$ on the line between R and the ideal attacker. Intuitively, with this geometry, the errors caused by the approximation are zero at R and small in the area nearby.

Receiver. The receiver estimates its unknown position P and clock bias δt by solving a system of at least four equations of the form:

$$\rho_j = ||P - S_j|| + c \cdot \delta t \tag{1}$$

Where ρ_j is the distance between the receiver and the satellite measured from the time of arrival of the signal, S_j is the satellite's position, and c is the speed of light. In practice, the receiver does not look at absolute values of ρ_j, but at the difference in arrival time across different satellites.

Single satellite. `GNSS-WASP` uses multiple transmitters, each impersonating one satellite. For simplicity, we first analyze the behavior of a single transmitter/satellite pair and then expand the results. Figure 5 depicts the attack geometry and timings. The goal of the attacker is to convince receivers near a reference position R that they are receiving signals from a satellite at position S_1 in the sky.[2] Ideally, any victim receiver V_i should receive the spoofing signal at a time proportional to its distance x from S_1. To approximate this behavior, the attacker places its transmitter at position A_1 on the line between the reference R and the satellite S_1, but at a lower altitude. The attacker delays its transmission time such that its signal arrives in R precisely at the same time it would arrive and was sent by S_1. In particular, the attacker transmits at $t_s + y/c$, where t_s is the transmission time that S_1 would use, y is the distance between A_1 and S_1, and c is the speed of light. The attacker signal arrives at R at $t_s + y/c + w/c$, where w is the remaining distance between A_1 and R, exactly like the signal from S_1 would. At a generic victim location V_i, the behavior is slightly different. While the signal from S_1 would start at t_s and travel through a distance x between S_i and V_i, the attack signal starts at $t_s + y/c$ and travels through a distance z between A_1 and V_i. On the positive side, the arrival time at V_i depends on its position as expected. On the negative side, the attack causes a timing error of $\Delta t = y/c + z/c - x/c$, corresponding to a distance error of:

$$\Delta d = y + z - x \tag{2}$$

The closer V_i to R and/or the closer A_1 to S_1, the smaller the geometrical and timing errors Δd and δt.

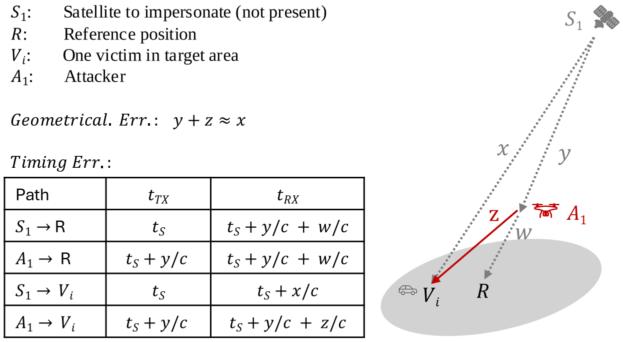

S_1: Satellite to impersonate (not present)
R: Reference position
V_i: One victim in target area
A_1: Attacker

$Geometrical.\ Err.:\ \ y + z \approx x$

$Timing\ Err.:$

Path	t_{TX}	t_{RX}
$S_1 \rightarrow$ R	t_S	$t_S + y/c\ +\ w/c$
$A_1 \rightarrow$ R	$t_S + y/c$	$t_S + y/c\ +\ w/c$
$S_1 \rightarrow V_i$	t_S	$t_S + x/c$
$A_1 \rightarrow V_i$	$t_S + y/c$	$t_S + y/c\ +\ z/c$

Figure 3: **Timing analysis for one satellite.** The attacker A_i impersonates satellite S_i in an area near the reference location R. For any victim V_i, the arrival time of the spoofing signal should be proportional to its distance x from S_i. To achieve this, A_i stays on the line between R and S_i, delaying its transmission proportionally to its distance y from S_i. R receives the signal exactly at the expected time. V_i receives it according to its position but with a small error caused by the imperfect geometry ($y + z \approx x$). The closer V_i is to R, the smaller the error. The closer A_1 is to S_1, the larger the area where the error is small. The overall positioning error depends on all the spoofed satellites the receiver uses.

Overall solution and error regions. During the attack, each range ρ_i estimated by the receiver is affected by an additional error term Δd_i per satellite S_1 caused by `GNSS-WASP`, calculated in eq 2. However, since the receiver uses a numerical solver, finding an analytical expression for the error on the position is not trivial. A better approach consists of running a numerical evaluation with and without attack. Figure 4 shows an example of such analysis on a 20 by 20 grid of receivers with 500 m spacing and attackers at 300 m from the reference. We set the clock bias to zero for simplicity and focus on the position error. We observe that the error does not grow uniformly, creating three main regions. The error is small near the reference location, and the consistency of absolute posi-

[2] The actual satellite 1 is either not visible at position R, or the attacker overpowers it. Only the spoofing signal from A_1 is taken by the receiver.

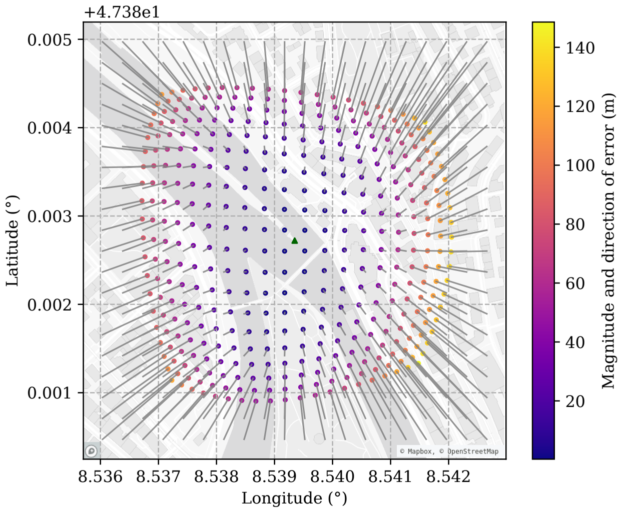

Figure 4: **Error analysis in a target area.** The attacker is placed at 300 m from the reference point (green triangle) at the center of the figure. Each of the dots represents the position reported by a receiver belonging to a 20 by 20 square grid over an area of 500 m by 500 m. The hue represents the magnitude of the error (including the error in altitude), while the grey lines connect the real position to the reported position (projected on the 2D map). The error is minimal in a large area but increases towards the edges. The relative error between any two points remains moderate.

tions and relative distances is well preserved. Further away, the absolute position error proliferates, but relative distances are still preserved well. Far from the reference, the satellites appear similar to a single source. As long as they are in range, victims believe to be located at the edge of the border area, and consistency is not preserved. Figure 5 shows the magnitude of the position error over different attacker distances from a given reference location and time.

These results also apply to real receivers and generalize to different locations, times, and attacker distances, as we discuss in detail in Sections 4 and 5. In particular, in Section 5, we study the error introduced by GNSS-WASP compared to legitimate noise and discuss possible countermeasures.

In Appendix A (Figure 15) we evaluate synchronization and positioning errors, and in Appendix B (Figure 16) we show trajectories of moving receivers on the error map.

3.5 GNSS-WASP Prototype

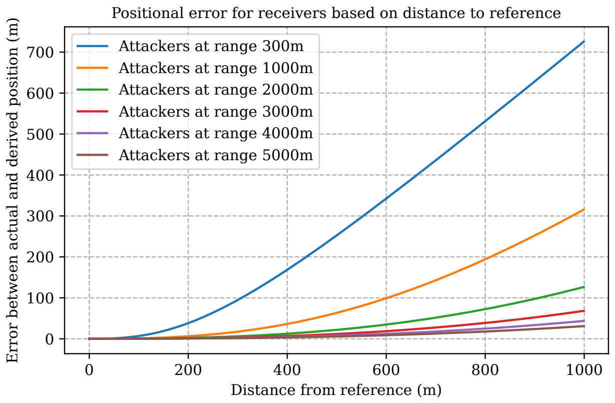

Figure 5: **Position error magnitude** estimated with the model. Curves are small and flat near the reference location (attack area), then start turning (border), and finally grow linearly (outside area). Greater satellite distance increases coverage.

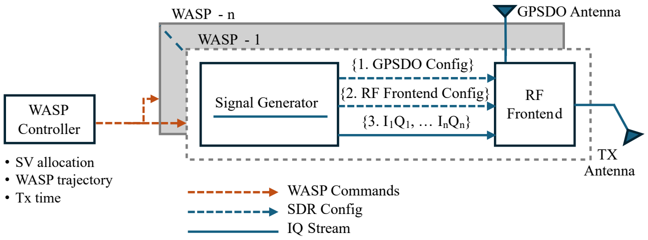

Figure 6: **Design Schematic** A WASP controller performs satellite allotment and sends a common transmission to WASP units. Each WASP unit is responsible for shadowing and impersonating a single satellite.

To demonstrate the feasibility GNSS-WASP and to evaluate its performance, we implement a prototype using commercially available off-the-shelf hardware based on the attack overview and the numerical model discussed in Section 3.3 and Section 3.4 respectively. Figure 6 presents a schematic representation of GNSS-WASP prototype. Two main components of GNSS-WASP are: i) the *WASP Controller* and ii) the *WASPs*. The WASP controller is responsible for satellite allocation and WASP synchronization. It also transforms satellite orbits into a flight path that a WASP can follow to shadow the assigned satellite (Refer to Figure 2). Each WASP is equipped with a signal generator that generates a signal as seen by the transmitter at its current location for the allocated satellite. It then interfaces with an RF frontend to transmit the generated samples precisely at the assigned time.

The success of GNSS-WASP depends on its ability to transmit satellite signals synchronously to avoid any inconsistencies in the pseudoranges derived by the receiver and to minimize the error explained in Section 3.4. Given the spatial diversity of WASPs, the accuracy requirements of satellite navigation, and the ease of availability and implementation, we use GPSDOs in GNSS-WASP for over-the-air clock synchronization. It combines a GPS receiver and a stable oscillator that uses a broadcast GPS signal as the timing source to discipline the oscillator and generate a stable 10 MHz signal, which an SDR can use to discipline its clock. The broadcast nature and wide-area coverage of GPS enable multiple SDRs

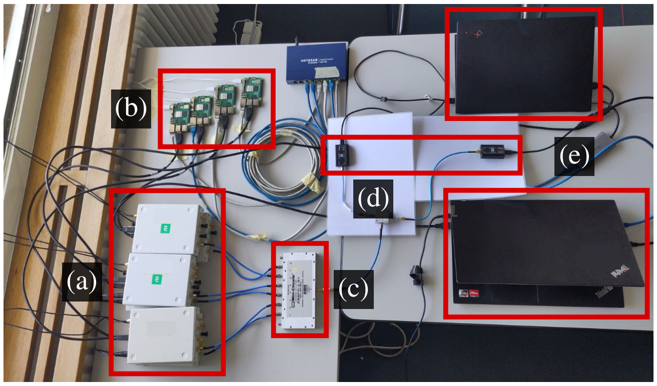

Figure 7: **GNSS-WASP Evaluation Setup (actual).** A photo of the actual prototype and the evaluation setup. Six USRP B210s equipped with GPSDOs **(a)** are controlled by four Raspberry Pis **(b)** and two consumer laptops **(e)**. A coordinator **(e)** interfaces with all the controllers over a local area network. It provides instructions on internal clock synchronization, spoofing signal generation, and transmission. These signals are combined by a splitter **(c)** and received by uBlox M9N receivers **(d)** for further offline evaluation.

to synchronize and have a common time and clock reference.

As per the specifications, with a proper GPS lock, Commercially available off-the-shelf (COTS) devices like USRP B210 using a GPSDO can achieve clock stability up to ±1 ppb [10]. Each WASP runs a real-time GPS signal generator derived from popular signal generation software GPS-SDR-SIM [24]. The signal generator interacts with the USRP through C UHD API [11] calls and configures the board to use GPSDO as a time and clock source. Next, we use the *has_time_spec* flag to set a common transmission time derived from GPS. It is important to note that GNSS-WASP is designed to be implemented with readily available hardware and software components.

4 Evaluation

This section presents a performance and security evaluation of GNSS-WASP. Specifically, we evaluate the stability and accuracy of the PVT solution estimated by hardware receivers manufactured by uBlox and a software receiver. Our primary performance evaluation metric is the error in the position estimated by the receiver and the target position represented in the form of distance.

4.1 Evaluation Setup

To validate the feasibility of our attack, we use COTS hardware and software for GNSS-WASP implementation. Our evaluation setup uses a consumer-grade laptop as the WASP controller. This controller connects to six WASPs over a local area network. We use USRP B210 software-defined radios [26],

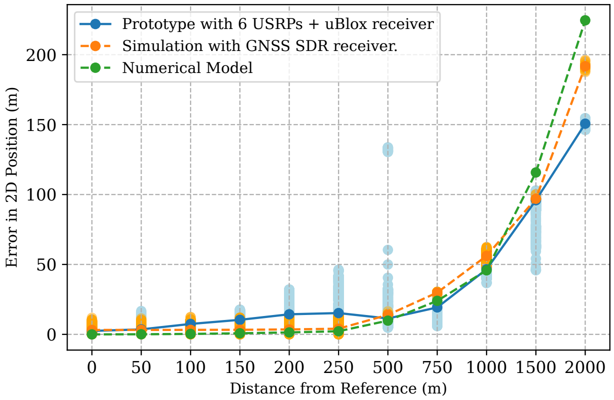

Figure 8: **GNSS-WASP Location Spoofing Accuracy** A comparison of location spoofing accuracy of the prototype.

each equipped with a GPSDO module to have a stable clock source and synchronize transmission time between all radios. Each USRP is connected to a Raspberry Pi 5 and is responsible for spoofing a single satellite.

Since we perform all our evaluations using a wired setup, we had to implement a simple channel model that adds appropriate signal delay and phase distortion as experienced by a receiver at a certain location. It is important to note that this is required because of the ethical, safety, and legal issues associated with over-the-air transmission of GNSS signals. We ensure that there is no signal leakage. This channel model is not required in a real attack scenario since the actual channel between the WASP and the receiver will add these. Finally, the transmitted signals are combined using Mini-Circuits ZB8PD-2-S+ splitter/combiner [5] and passed on to a GNSS receiver. We evaluate our attacks on two different receiver models, uBlox XPLR M9 and uBlox MAX M8Q. Figure 7 shows an image of the actual evaluation setup.

4.2 Location Accuracy and Stability

As described earlier in Section 3.4, the victim receiver will experience a numerical error the farther it is from the reference position. We used our prototype to evaluate how a real receiver processes the spoofed signal generated by USRPs synchronized with GPSDOs. In this test, we select a reference position in Paris and generate eleven scenarios where the spoofed location is at a specific distance from the reference position. Since we are limited to wired transmissions, this approach helps to evaluate the performance of GNSS-WASP at different locations. It is important to note that we must simulate the channel in our evaluation setup since we are not transmitting the generated signal over the air; however, this step is not required in real life. The test sequence is as follows: first, we use the WASP controller to set a common transmission time once all the USRPs are locked to the GPS

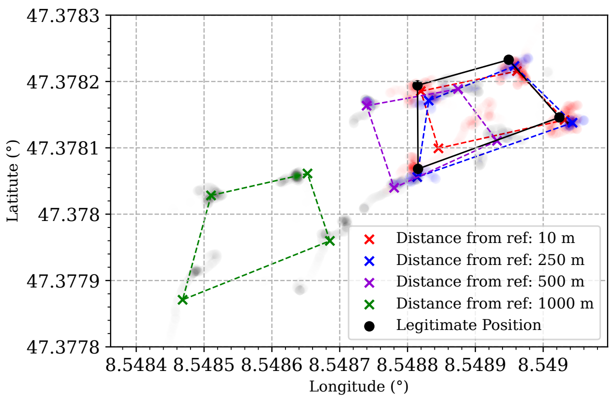

Figure 9: **Four receivers with known separation.** We place four receivers at known distances and collect legitimate position data overnight. Then, we connect them to GNSS-WASP and test four spoofing scenarios at different distances from the reference position. Both the average positions (black dots and colored markers) and the distribution of single measurements (shadowed dots) are visible. The figure intuitively shows that i) even legitimate measurements are affected by noise, ii) the absolute position error (which cannot be calculated by a victim in lack of ground truth) increases with the distance from the reference, but iii) within the uncertainty caused by noise the relative distances across receivers remain consistent, preventing detection with classic multi-receiver countermeasures that would expect all positions to collapse in one point under a single-antenna attack. Detailed numerical data are shown in Table 2 and an evaluation of GNSS-WASP-aware countermeasures is presented in Figure 13.

signal. Second, we configure the channel model to simulate the channel experienced at the target location when a WASP is 5 km from the reference position. Third, we start the uBlox receiver and dump geo-coordinates. We repeat this for every target location. Figure 8 shows the results of this experiment, where we plot the error in the received and target locations. Along with data gathered from real receivers, we also include data from an offline simulation that uses GNSS-SDR and the values from the numerical model for comparison. Here, we can observe that a real receiver follows the error estimation obtained through numerical modeling from Section 3.4. However, the two traces contain some irregularities, as the real receivers are black boxes, which are generally considered non-deterministic. Moreover, during each run, GPSDOs may drift depending on environmental conditions.

Figure 8 shows the results of this experiment, where we plot the error in the received and target locations. Along with data gathered from real receivers, we also include data from an offline simulation that uses GNSS-SDR and the values from the numerical model for comparison. Here, we can observe that a real receiver follows the error estimation obtained through numerical modeling from Section 3.4. However, the two traces contain some irregularities as the real receivers are black boxes and are generally considered non-deterministic. Moreover, during each run, GPSDOs may drift depending on environmental conditions.

4.3 Attack on Multi-receiver Setup

One of the main benefits of GNSS-WASP over conventional spoofing techniques is that GNSS-WASP is agnostic of receiver location, and it can spoof multiple receivers without experiencing the group-spoofing problem described in [62]. Assuming a conventional single antenna attacker model, several works like [34,68] have proposed countermeasures that use multiple receivers to identify a spoofing attack. To show resilience to such a strategy, we performed an experiment comparing the error distribution of a formation of four hardware receivers processing legitimate signals and processing signals generated by GNSS-WASP. A formation of four receivers was also used in [34] to successfully evaluation of the multi-receiver countermeasure against spoofing attacks. For this test, we first set up four Raspberry PIs with a Uputronics Raspberry Pi+ GPS Expansion Board running a Ublox MAX-M8Q GPS chip on the rooftop of our building. We let them gather data overnight. Next, we create four attack scenarios where we set the reference point at a distance of 10 m, 250 m, 500 m, and 1000 m. The position of our attack transmitters is decided only by the position of the reference and not by the individual positions of the receivers under attack. We follow the same methodology as Section 4.2 to gather data. In total, we ran sixteen scenarios (four for each receiver) and evaluated the error distribution over a 5-minute trace. Results are shown in Figure 9 and Table 2. In Section 5.4 we will evaluate GNSS-WASP-aware countermeasures on these data.

5 WASP-Aware Countermeasures

Countermeasures based on multiple receivers [34] and sensor-fusion [52] work under the assumption that practical attacks (e.g., single transmitter) cause large inconsistencies and advanced strategies (e.g., multiple transmitters) are impractical. In their original form, they are ineffective at detecting GNSS-WASP within the 3D area where consistency is preserved. In this section, we study to which extent GNSS-WASP can be identified from the approximation error it introduces, and we discuss strategies to improve detection.

5.1 Detection Game

In line with previous work on countermeasures, through a potential victim's perspective, we informally define a security game involving the legitimate constellation, the attacker who wants to impersonate it, and a defender with grids of moving receivers placed in known formations. The goal of the defender is to reliably distinguish whether its position was derived from the legitimate satellites or from the attacker impersonating them. More formally, the defender attempts to distinguish the following two hypotheses:

- H_0: The receivers are not under spoofing and are computing their position using legitimate signals from legitimate satellites.
- H_1: The receivers are under spoofing of a `GNSS-WASP` attacker and are computing their position from the signals transmitted by the attacker.

In practice, neither the attacker nor the defender always win. On the one hand, the test occurs under noisy conditions, and the defender has some probability of making wrong decisions. On the other hand, `GNSS-WASP` introduces a small but deterministic position error that helps detection. The defender is characterized by the tradeoff between False Acceptance Rate (FAR) and False Rejection Rate (FRR):

- $FAR = P(Decision = H_0 \mid H_1)$: the probability of accepting spoofed measurements as legitimate.
- $FRR = P(Decision = H_1 \mid H_0)$: the probability of rejecting legitimate measurements as spoofed.

As GNSS is typically used in real-time applications such as for the navigation of moving vehicles, having low decision latency and low false rejection rates is generally important.

5.2 Legitimate-vs-Attacker Error Modeling

We model the position estimation with and without spoofing in place. In both cases, the receiver finds its position P and clock error δt by solving a system of equations of the form:

$$\rho_i = ||P - S_i|| + c \cdot \delta t \tag{3}$$

where the *pseudorange* ρ_i is the noisy estimate of the distance between satellite and receiver obtained from the arrival time, S_i is the position of the legitimate satellite, $||P - S_i||$ is their geometrical distance, and c is the speed of light. In the legitimate case, the pseudorange can be written as:

$$H_0: \ \rho_i^{H_0} = ||P - S_i|| + c \cdot \delta t + \underbrace{\delta d}_{noise} \tag{4}$$

where $c \cdot \delta t$ is the error caused by the receiver's clock being less accurate than the satellite's one, and δd is the measurement error due to other factors such as the receiver's noise, multipath, and ionospheric and tropospheric delays. Previous work has shown that the variance of this error can amount to several meters [2]. Like the position P, also the clock error δt is an unknown that the receiver estimates by solving at least four equations. Under a `GNSS-WASP` attack, the pseudorange is:

$$H_1: \ \rho_i^{H_1} = ||P - S_i|| + c \cdot \delta t + \underbrace{\delta d}_{noise} + \underbrace{\Delta d_i}_{attack\ bias} \tag{5}$$

where the additional term Δd_i is the deterministic error introduced by the `GNSS-WASP` attack geometry. This error can be calculated from the positions of the attacker, satellite to impersonate, victim and reference, as we explained in Section 3 and Figure 3. The attacker advertises its position as that of the satellite to impersonate, hence the term S_i remains the same even if the attacker is located at much lower altitude.

The main advantage for the defender, and problem for the attacker, is that the geometry of the `GNSS-WASP` attack introduces a deterministic bias Δd_i on the pseudorange estimates. Even if it is minimal within the 3D area, it can eventually be detected by a motivated defender either with a sufficient number of measurements or by exiting the area. For simplicity, we assume that the noise δd affects both legitimate satellites and attacker transmitters in a similar way. This is because receiver noise affects the reception of both, and other effects can be generally predicted and simulated by the attacker.

As explained in Section 3, legitimate noise and attacker bias propagate to the position estimate through the numerical solution of the system of equations performed by the receiver, and can be evaluated numerically. The positions become:

$$H_0 : P^{H_0} = P + \underbrace{\delta P}_{noise} \tag{6}$$

$$H_1 : P^{H_1} = P + \underbrace{\delta P}_{noise} + \underbrace{\Delta P}_{attack\ bias} \tag{7}$$

where δP is the legitimate noise caused by δd_i and ΔP is the bias caused by Δd_i of each satellite. Figure 5 and Figure 4 in Section 3 present examples of numerical evaluation of ΔP. The error might change over time as the constellation moves. A defender knowing the distance between two receivers can try to distinguish the following two hypothesis:

$$H_0 : d^{H_0} = ||P_A - P_B + \underbrace{\delta P_A - \delta P_B}_{noise}|| \tag{8}$$

$$H_1 : d^{H_1} = ||P_A - P_B + \underbrace{\delta P_A - \delta P_B}_{noise} + \underbrace{\Delta P_A - \Delta P_B}_{attack\ bias}|| \tag{9}$$

where $\Delta P_A - \Delta P_B$ is the combined effect of the bias at A and B. This case also covers a defender with one receiver that moves from A to B and knows the distance it traveled (e.g., using inertial sensors). As the errors in nearby locations have similar directions and magnitudes, the error on the relative distance is often small even if the absolute error on the individual positions is large, increasing the attack area.

The defender uses a statistical test to decide whether a given position is legitimate (H_0) or spoofed (H_1). For simplicity, we use a similar approach as [34], but adapted to our specific models for H_0 and H_1, as follows. The defender estimates the distribution of distances for H_0 under legitimate conditions, and chooses a threshold γ to detect outliers. The threshold can be tuned to obtain different tradeoffs between FAR and FRR. When evaluating grids with more than 2 receivers, we measure all possible distance pairs and declare spoofing if one of them is an outlier. Indeed, the direction of the attacker bias changes with the receiver location (see Figure 4 in Section 3) and might match one of the pairs but not the others.

Difference with prior work. Consider two receivers at positions P_A and P_B, with distance $||P_A - P_B|| = 50\,\text{m}$, and $5\,\text{m}$ of legitimate error on each. In the legitimate scenario the distance is $d^{H_0} = (50 \pm 10)\,\text{m}$, whereas under an attack with single transmitter both receivers compute the same position and the distance is $d^{H_1} = (0 \pm 10)\,\text{m}$. The defender can easily distinguish the two cases [34]. Alternatively, assume that the two positions are within the area of a GNSS-WASP attack, and that the bias is $\Delta P_A = 2\,\text{m}$ and $\Delta P_A = -2\,\text{m}$. The distance under attack becomes $d^{H_1} = (54 \pm 10)\,\text{m}$, and cannot be easily distinguished from the legitimate d^{H_0} anymore. GNSS-WASP might either increase or decrease the distance. Hence outliers can be either larger or smaller than the legitimate distribution.

Advantages and Limitations of Repeated Measurements. A countermeasure could leverage multiple independent measurements to distinguish the attack bias from Gaussian noise. We test two different methods: aggregating measurements over a sliding window and CUSUM [30, Section 6.3.2.3]. While aggregating measurements improves detection, it also has limitations. First, it increases the detection latency, which might be problematic for vehicles traveling at high speed. Second, the bias is deterministic for a given time and location but varies over time with satellites' movement and across positions with victims' movement. Hence, its effect across multiple measurements does not necessarily always add-up as it can also decrease (e.g., see experimental data in Figure 17 in Appendix B). Finally, biases in legitimate measurements (e.g., due to multi-path in urban areas) can cause false positives.

Generality. Our model describes the attack bias on absolute and relative positions, and hence the fundamental reason why detection is possible regardless of the method applied. Only relative errors are useful, as the defender does not know its absolute position, but can compare differential measurements (e.g., distances, acceleration). In Section 7 we provide a broader overview of existing countermeasures.

5.3 Numerical Evaluation

To evaluate the countermeasures, we use the mathematical model that we have presented in Section 3, to which we add the capability of adding noise on the pseudoranges. In favor of the defender, we consider only the noise caused by the receiver (accounting for around 0.3 m to 1 meter of deviation) [2]. We exclude other sources of error, such as ionospheric delay and multi-path (accounting for several meters of error) whose effect can be observed in Figure 9. Typical consumer receivers estimate their position at 10 Hz

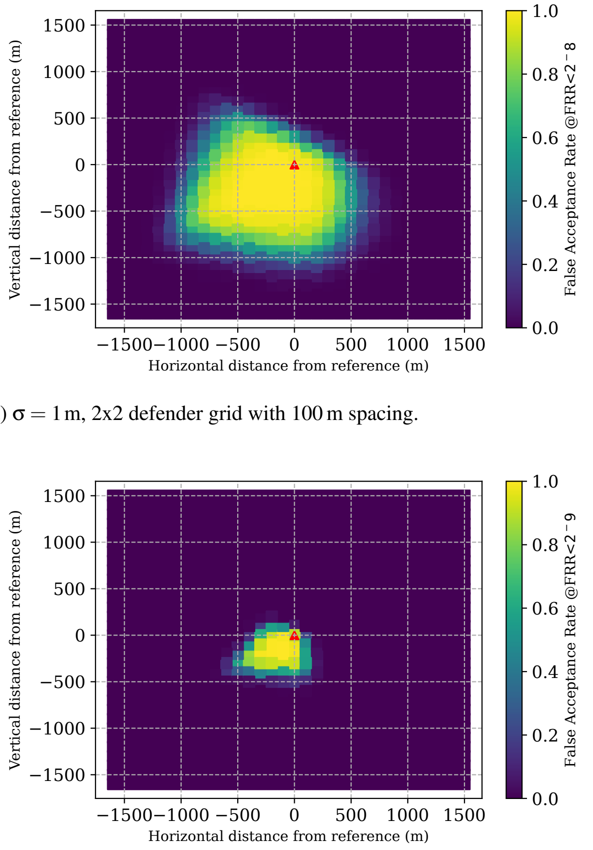

(a) $\sigma = 1\,\text{m}$, 2x2 defender grid with 100 m spacing.

(b) $\sigma = 0.3\,\text{m}$, 2x2 defender grid with 100 m spacing.

Figure 10: **False acceptance rate.** The maps show the FAR for GNSS-WASP in a $10.24\,\text{km}^2$ area around the reference (red marker). The defender has a grid made of 4 receivers spaced by 100 m and the attackers are at 5 km from the reference. The FRR is 2^{-9}. We vary the level of receiver noise between 1 m (a) and 0.3 m (b) and we neglect other sources of legitimate errors such as multi-path that would favor the attacker. Despite the favorable conditions for the defender, GNSS-WASP is successful in a large area (yellow) in both cases: more than $1\,\text{km}^2$ for (a) and approximately $0.25\,\text{km}^2$ for (b).

Moving Receiver Grids. We investigate the effectiveness of receiver grids to detect GNSS-WASP with 8 attackers placed at 5 km from the reference. The receivers are placed at known relative distances from each other (e.g., on a large vehicle or on the members of a swarm). Since the grids move over time, the ground truth position is assumed to be unknown.

We consider a large 2 by 2 grid made of 4 receivers with a spacing of 100 m. The grid measures the known distances on each of the 6 non-directed edges between nodes. Given the large spacing and multiple edges in different directions, this scenario favors the defender. It would be practical for large vehicles (e.g., boats) but not for smaller ones (e.g., cars, drones) unless they travel in a formation. For the evaluation, we con-

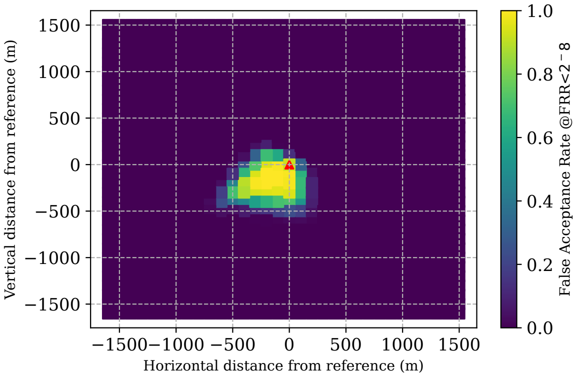

Figure 11: **False acceptance rate with aggregation**. Aggregating measurements over a 1 s rolling window reduces the FAR compared to Figure 10a. However, such latency might be already unacceptable at high speeds.

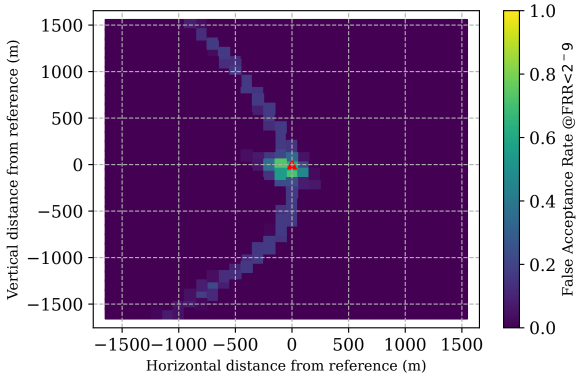

Figure 12: **False acceptance rate with CUSUM**. Detection is even more effective than with a sliding window.

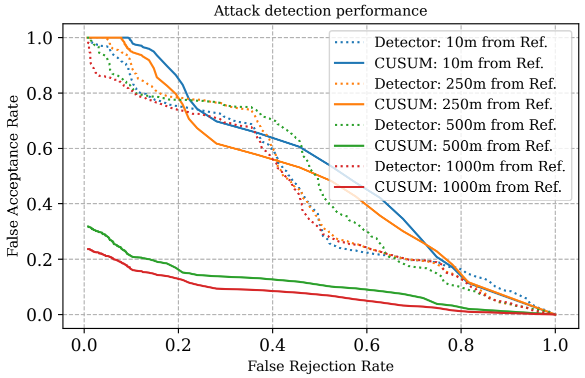

Figure 13: **Countermeasure on real hardware**. FAR-FRR curve for different distances from the reference. The FAR is high for small FRR, and lower for larger distance. CUSUM is generally better.

sider a grid of regularly spaced locations in a 10.24 km^2 area around a given reference.

We first estimate the legitimate noise by collecting 10000 measurements at each of 100 locations randomly selected in the area. We compute the distances between the defender's nodes using all three dimensions (longitude, latitude, and altitude). We normalize the distance measured on each edge by subtracting its known length. Based on these data, we select a threshold such that the False Rejection Rate is 2^{-9}. Such large FRR favors the defender, as smaller values would require more generous thresholds and increase the attacks' success rate. We do not assume a specific data distribution and simply run the countermeasure for multiple threshold values until the correct one is found.

We test our countermeasure on a different day, collecting $10,000$ ideal measurements and $10,000$ measurements with GNSS-WASP for each possible defender's grid location. Set #1 corresponds to the case when the receivers are receiving their signals from the legitimate satellites (H_0). In contrast, set #2 corresponds to the receivers being under spoofing from the GNSS-WASP transmitters. We then estimate the False Acceptance Rate as the number of measurements for which the grid does not detect the attack over the total number of measurements in the attack scenario and the FRR as the number of false alarms over the total number of measurements in the legitimate scenario, assuming that the defender uses the threshold identified in the previous phase.

We repeat the experiments for two different levels of receiver noise on the pseudoranges: standard deviation $\sigma = 0.3$ m and $\sigma = 1$ m. Recall that these cases *favor* the defender. We also conduct a similar evaluation in Section 5.4, comparing the attack performance with data collected in the field.

Results are shown in Figure 10. On the one hand, they demonstrate that GNSS-WASP is successful on large areas despite the countermeasure. On the other hand, they highlight that the approximation error introduced by GNSS-WASP can eventually be detected by a grid moving away from the reference (unless the attacker updates the reference).

Evaluating repeated measurements. We evaluate two detection strategies; i) moving average of 10 samples, and ii) CUSUM, with a detection sigma of 1. In both cases, we train the thresholds and test the detector on the same datasets as before. Results are shown in Figures 11 and 12.

5.4 Evaluation with Real Hardware

We also analyze the countermeasure for the attack on a multi-receiver setup that we have presented in Section 4, shown in Figure 9. We use one set of real measurements taken from the two commercial receivers to estimate the threshold that gives a desired FRR. Then we test the FRR using our prototype (connected via cable). Figure 13 shows the FRR as a function of the desired FAR, for the attack at three different distances from the reference. For low FRR the FAR is high.

6 Discussion

Limitations: The most influential factor that ensures the success of GNSS-WASP is precise time synchronization between all the radios. Satellite navigation heavily depends on precise clocks and is extremely sensitive to clock drifts, and thus, GNSS-WASP can easily break if the clocks are not synchronized. The GPSDOs we use in our prototype guarantee synchronization errors of ±50 *ns* with clock stability up to ±1 ppb. The ideal operating temperature is 25°C. During our experimental evaluation, we observed that factors like ambient temperature and uniformity in GPSDO antenna cable lengths significantly impact the stability of the clocks, thus causing the USRPs to be out of sync. As seen in Figure 14, these effects directly impact the position calculated by the receiver. Thus, we encountered difficulties in properly evaluating GNSS-WASP on a warm summer day. However, an attacker can send over-the-air clock corrections to WASPs like the real GNSS by setting up monitoring stations at known locations. Our attack strategy involves high-flying drones that shadow real satellites. In practice, it is possible that these drones do not track the satellite's trajectory perfectly. Even though drones capable of achieving centimeter precision with RTK-GPS are available nowadays [6, 8], they are still sensitive to wind, thus introducing another source of error. Since the coverage of GNSS-WASP is limited, a swarm of dedicated receivers can move around strategically to detect the attack and even find a path out of the spoofing region. However, extensive work is being done on identifying and tracking drones and other vehicles through audio localization [16, 67], RF localization [18, 60], and even optical and thermal sensors [21, 58]. Integration of these technologies will enable GNSS-WASP to roughly follow the swarm's movement to keep it inside the target area. Unlike legitimate satellite navigation systems with global coverage, GNSS-WASP and other spoofing attacks have limited coverage. Any target transitioning between the affected and non-affected areas experiences a jump in its position. Its magnitude depends on several factors like location offset introduced by GNSS-WASP and the direction in which a target enters the spoofing region. Depending on the magnitude of this jump, a receiver may raise an alarm by detecting high variance in the PVT solution or inconsistency in INS/GNSS fusion. Notably, this jump will be an isolated incident because it will occur only when the target transitions; once the target has acquired the signal from GNSS-WASP, the receiver will not experience these jumps. Moreover, given the inherent noise in GNSS, most countermeasures observe a window of several measurements to avoid false positives.

Gap to Real-World Demonstration: We have demonstrated the feasibility of GNSS-WASP with a mathematical model and a proof-of-concept implementation with real hardware over a wired connection. A full demo would require over-the-air transmission of real GPS signals, which would violate regulations and might cause harm. It would also require additional engineering efforts to deploy the RF equipment on the drones, and design appropriate flight control systems. The altitudes, speeds, and transmission powers required are within reasonable values for the capabilities of existing technology. Executing the attack may involve additional techniques often combined with spoofing, such as jamming receivers to force a cold start or gradually overtaking victims by smoothly adjusting power and position. The attack performance would depend on the quality of the drones, and external weather conditions such as wind speed and temperature, and would require an extensive field evaluation in different conditions.

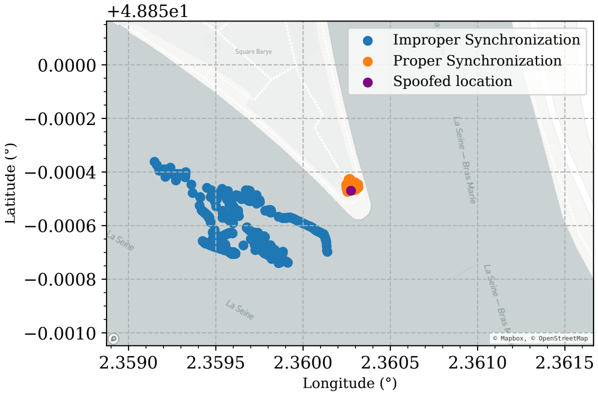

Figure 14: **Effect of GPSDO alignment.** Synchronization amongst transmitters depends on the stability of GPSDO. When clocks are synchronized (orange trace), the spoofed location is very close to the target location with an average error of only 2.0768 m. In contrast, unsynchronized clocks (blue trace) result in a larger average error of 48.9977 m.

Extending Selective Delay Attacks: Owing to the broadcast nature of GNSS signals, it is possible to launch relay/replay attacks where an attacker records GNSS signals in one place and then either relays the signals near the target or replays the signal on a later date. These practical attacks have been demonstrated in works like [41] and [36]. The main advantage of these attacks is that they are effective even against cryptographic countermeasures, such as the recently launched OSNMA. GNSS-WASP can be extended to support a selective delay attack using directional antennas and relaying satellite signals with a slight delay. This approach is inspired by Motallebighomi et al. [41]. However, it will relay the original signal with an added delay instead of re-generating it.

Minimizing The Effect of Bias: The bias introduced by GNSS-WASP can be reduced by increasing the altitude of the transmitters, but cannot be eliminated unless they are placed in orbit. As the error is smaller close to the reference, the attacker can also coarsely update the reference to keep it close to the group of receivers it wants to target, still without having to finely track and spoof each individual receiver.

7 Related Work

Satellite navigation plays a crucial role in modern systems. It serves as a source of accurate position information and precise timing information. The importance of satellite navigation systems in safety—and security-critical applications and their fundamentally insecure design have piqued the interest of authorities, academicians, and some antisocial elements. The earliest accounts of the susceptibility of GPS to spoofing and jamming date back to 1992 [64]; since then, our understanding of GNSS threats and spoofing capabilities has come a long way. The comprehensive analysis presented by Tippenhauer et al [62] enhanced the understanding of GPS spoofing. They state the requirements for a successful GPS spoofing attack and present countermeasures that leverage the spatial limitations of a single antenna attacker. Since then, several works have been published that focus on demonstrating the impact of GNSS spoofing and jamming [40]. These include power grids [57], road navigation systems [19, 43, 69], autonomous vehicles [35,47,54] and yachts [63]. The scientific community is also engaged in active research on countermeasures. On a high level, they can be classified as follows: **1. Physical layer anomaly detection** techniques that monitor various parameters like noise level, clock errors, the angle-of-arrival of the signal, and cross-ambiguity function. For instance, in [51, 66], authors present a strategy to detect auxiliary peaks to identify malicious signals. In [17, 37] authors provide a spoofing detection mechanism based on identifying errors in angle-of-arrival and direction of arrival of satellite signals. These methods generally apply to any attack including GNSS-WASP, though GNSS-WASP has an advantage: the physical-layer characteristics of GNSS-WASP's signals (e.g., delay, power, angle of arrival) naturally change in different areas of the target area, consistently with the distance and angle from each individual transmitter (and thus with the satellite that it impersonates). **2. Cryptographic countermeasures**, this category of countermeasure focuses on providing authenticity and confidentiality of navigation messages through using cryptographic primitives. Since GNSS is designed to be a free and open-to-all system, keeping the spreading codes secret is not feasible. Broadly, there are two strategies for verifying the authenticity of received signals: i) navigation message authentication (NMA) [29, 65] or ii) spreading code authentication (SCA) [13]. As discussed in Section 6, authentication would complicate the attack implementation (e.g., requiring a relay) without fundamentally preventing it. **3. PVT Consistency**, these countermeasures focus on spotting inconsistencies in the derived PVT solution across multiple systems, including other satellite navigation systems [46], crowd-sourcing [33], and correlation between civilian and military signals [49]. As long as a victim can obtain another trusted positioning source with sufficient accuracy, even an ideal attack impersonating the full constellation would be detected from the inconsistent spoofed position. **4. Sensor-fusion**, several works like [22,38,52,61] leverage the short term accuracy of inertial sensors to detect errors in GPS measurements through sensor-fusion algorithms. These works leverage the sensitivity of sensor-fusion algorithms to inconsistencies in the measurements. Similar observations as for countermeasures based on multiple receivers apply. Assuming that a receiver starts a trajectory from within the target area, its measured position would change consistently with its movements, and thus with other sensors like INS. Only the attack bias as discussed in Section 5 could eventually reveal the attack. However, while the uncertainty on the fixed position between two receivers can be very low, the output of INS might experience significant drift over time and needs to be corrected with GPS itself.

Despite the development of effective countermeasures, including cryptographic countermeasures and robust receivers, GNSS technology still remains vulnerable to physical layer relay/replay attacks as shown in [36, 41] and forward-error estimation attack demonstrated in [48]. Even countermeasures that rely on combining authenticated and non-authenticated signals can be defeated, as demonstrated in [14]. Moreover, most of these works, including both attacks and countermeasures, focus on the naive single-antenna attacker and fail to consider a more complex multi-antenna attacker like GNSS-WASP.

8 Conclusion

Global Navigation Satellite Systems provide accurate positioning for millions of users and many sensitive applications. In this paper, we have proposed GNSS-WASP, a novel spoofing attack that can simultaneously and consistently manipulate the position of many victim receivers in a target area without any prior knowledge or tracking of their position. This enables novel attack scenarios such as diverting entire swarms of moving vehicles with minimal effort. The error introduced by the attack is minimal and it is hidden in the natural noise and inaccuracy of the measurements. We have evaluated GNSS-WASP by means of extensive modeling, simulation, and testing on real GPS receivers with a prototype implementation based on off-the-shelf software-defined radios. GNSS-WASP can spoof two receivers while maintaining their relative distance with an average error of 0.97 m for locations 1000 m away from the reference position. We have shown that current state-of-the-art countermeasures based on checking the consistency of positioning across multiple receivers or sudden movements can be bypassed with high success rate by GNSS-WASP, and we have proposed improved detection methods.

9 Acknowledgement

This research has received funding from the Swiss National Science Foundation under NCCR Automation, grant agreement 51NF40_180545. Map data are from OpenStreetMap using Mapbox.

10 Ethics Considerations and Compliance With the Open Science Policy

The study presented in this paper did not involve any human subjects. Following the best practices in this domain, the experiments involving transmission of GNSS signals were carried over wired connections and at low power in a controlled environment, ensuring the absence of any leak to the wireless spectrum and preventing any interference with the real system. After presenting a novel spoofing attack, we thoroughly developed and evaluated better defense strategies to mitigate the risk.

To facilitate the replication of our findings and in accordance with the open science policy, we published the Python code and configurations that we used to model the effects of GNSS-WASP on receivers and submit it for artifact evaluation. More specifically, we included all Python scripts required for: (a) data generation and visualization for error magnitude on a map in Figures 4, 16, (b) evaluation of error magnitude over distance from reference/distance of attacker in Figure 5, (c) modeling of attack detection for multiple receivers over time, with noise on pseudo-ranges in Figures 10, 11, 12, (d) modeling of position and synchronization errors in implementing the attack in Figure 15. We also released the datasets and analysis code required to generate: (e) the experimental spoofing accuracy in Figure 8, (f) the experiment with a 4-by-4 receiver grid in Figures 9, 13, 17, and Table 2, (g) the experimental synchronization error in Figure 14.

The export of dual-use technology is strictly regulated by export control laws worldwide. This is receiving increasing attention, including at our institution [25], and requires careful assessment. Even though some of the components of our prototype are open-source [24] and implement known attacks, our work presents a new attack with stronger capabilities. Due to potential dual use and export control considerations, and after consulting with the export control team at our institution, we decided not to release the code used to setup and run the attack on physical hardware (shown in Figures 6 and 7), used to generate the datasets for Figures 8, 9, 13, 14, and Table 2.

Artifacts available at https://zenodo.org/records/14734238.

References

[1] ADAM Pluto. https://www.analog.com/en/resources/evaluation-hardware-and-software/evaluation-boards-kits/adalm-pluto.html.

[2] GNSS Error Budget. https://www.vectornav.com/resources/inertial-navigation-primer/specifications--and--error-budgets/specs-gnsserrorbudget.

[3] GNSS Simulatio0ns. https://www.rohde-schwarz.com/ch-en/products/test-and-measurement/digital-standards/gnss-simulation_63493-1124871.html.

[4] HackRF One. https://greatscottgadgets.com/hackrf/one/.

[5] Mini-Circuits ZB8PD-2-S+ Splitter/Combiner. https://minicircuits.com/WebStore/dashboard.html?model=ZB8PD-2-S%2B.

[6] Phantom 4 RTK. https://enterprise.dji.com/phantom-4-rtk.

[7] PNT Simulation Systems. https://www.spirent.com/products/pnt-simulation-systems.

[8] RTK GPS Correction (Fixed Baseline). https://ardupilot.org/copter/docs/common-rtk-correction.html.

[9] Simulators. https://www.navtechgps.com/departments/simulators/.

[10] USRP B200/B210 Product Overview. https://www.ettus.com/wp-content/uploads/2019/01/b200-b210_spec_sheet.pdf.

[11] USRP Hardware Driver and USRP Manual . https://files.ettus.com/manual/page_c_api.html.

[12] Drone Swarm Technologies. 2023. https://www.gao.gov/products/gao-23-106930.

[13] Jon M Anderson, Katherine L Carroll, Nathan P DeVilbiss, James T Gillis, Joanna C Hinks, Brady W O'Hanlon, Joseph J Rushanan, Logan Scott, and Renee A Yazdi. Chips-message robust authentication (Chimera) for GPS civilian signals. In *Proceedings of the 30th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2017)*, pages 2388–2416, 2017.

[14] Francesco Ardizzon, Laura Crosara, Stefano Tomasin, and Nicola Laurenti. On mixing authenticated and non-authenticated signals against gnss spoofing. *IEEE Transactions on Information Forensics and Security*, 2024.

[15] ardusimple. Enable 1-centimeter precision on your drone. https://www.ardusimple.com/enable-1-centimeter-precision-on-your-drone/. Accessed 2025-01-11.

[16] Andrea Bernardini, Federica Mangiatordi, Emiliano Pallotti, and Licia Capodiferro. Drone detection by acoustic signature identification. *electronic imaging*, 29:60–64, 2017.

[17] Sriramya Bhamidipati, Kyeong Jin Kim, Hongbo Sun, and Philip V Orlik. GPS spoofing detection and mitigation in PMUs using distributed multiple directional antennas. In *ICC 2019-2019 IEEE International Conference on Communications (ICC)*, 2019.

[18] Udita Bhattacherjee, Ender Ozturk, Ozgur Ozdemir, Ismail Guvenc, Mihail L Sichitiu, and Huaiyu Dai. Experimental study of outdoor UAV localization and tracking using passive RF sensing. In *Proceedings of the 15th ACM Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization*, pages 31–38, 2022.

[19] James V Carroll, Karen Van Dyke, John H Kraemer, and Charles Rodgers. Vulnerability assessment of the US Transportation infrastructure that relies on GPS. In *Proceedings of the 14th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 2001)*, pages 975–981, 2001.

[20] Marco Ceccato, Francesco Formaggio, and Stefano Tomasin. Spatial GNSS spoofing against drone swarms with multiple antennas and wiener filter. *IEEE Trans. Signal Process.*, 68:5782–5794, 2020.

[21] Frank Christnacher, Sébastien Hengy, Martin Laurenzis, Alexis Matwyschuk, Pierre Naz, Stéphane Schertzer, and Gwenael Schmitt. Optical and acoustical uav detection. In *Electro-Optical Remote Sensing X*, volume 9988, pages 83–95. SPIE, 2016.

[22] Sagar Dasgupta, Mizanur Rahman, Mhafuzul Islam, and Mashrur Chowdhury. A sensor fusion-based GNSS spoofing attack detection framework for autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 23(12):23559–23572, 2022.

[23] DJI. D-RTK 2 high precision GNSS mobile station. https://www.dji.com/ch/d-rtk-2. Accessed 2025-01-11.

[24] Takuji Ebinuma. GitHub - osqzss/gps-sdr-sim: Software-Defined GPS Signal Simulator. https://github.com/osqzss/gps-sdr-sim.

[25] ETH Zurich. ETH Zürich export control for software, 2025. https://ethz.ch/en/industry/researchers/licensing-software/software-export-control.html.

[26] a National Instruments Brand Ettus Research. USRP B210 USB Software Defined Radio (SDR) - ETTUS Research.

[27] European Unition Aviation Safety Agency. EASA updates SIB on GNSS Outage and Alterations. https://www.easa.europa.eu/en/newsroom-and-events/news/easa-updates-sib-gnss-outage-and-alterations. Accessed 2024-07-31.

[28] Yanming Feng, Jinling Wang, et al. GPS RTK performance characteristics and analysis. *Positioning*, 1(13), 2008.

[29] Ignacio Fernández-Hernández, Vincent Rijmen, Gonzalo Seco-Granados, Javier Simon, Irma Rodríguez, and J David Calle. A navigation message authentication proposal for the Galileo open service. *NAVIGATION: Journal of the Institute of Navigation*, 63(1):85–102, 2016.

[30] N. Heckert, James Filliben, C Croarkin, B Hembree, William Guthrie, P Tobias, and J Prinz. Handbook 151: Nist/sematech e-handbook of statistical methods, 2002-11-01 00:11:00 2002.

[31] Todd Humphreys. Statement on the vulnerability of civil unmanned aerial vehicles and other systems to civil gps spoofing. *University of Texas at Austin (July 18, 2012)*, pages 1–16, 2012.

[32] Todd Humphreys, B. Ledvina, Mark Psiaki, B. O'Hanlon, and Jr Kintner. Assessing the Spoofing Threat: Development of a Portable GPS Civilian Spoofer. pages 2314–2325, 01 2008.

[33] Kai Jansen, Matthias Schäfer, Daniel Moser, Vincent Lenders, Christina Pöpper, and Jens Schmitt. Crowd-GPS-Sec: Leveraging crowdsourcing to detect and localize GPS spoofing attacks. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 1018–1031. IEEE, 2018.

[34] Kai Jansen, Nils Ole Tippenhauer, and Christina Pöpper. Multi-receiver GPS spoofing detection: Error models and realization. In *Proceedings of the 32nd Annual Conference on Computer Security Applications*, pages 237–250, 2016.

[35] Andrew J Kerns, Daniel P Shepard, Jahshan A Bhatti, and Todd E Humphreys. Unmanned aircraft capture and control via GPS spoofing. *Journal of field robotics*, 31(4):617–636, 2014.

[36] Malte Lenhart, Marco Spanghero, and Panagiotis Papadimitratos. Relay/replay attacks on GNSS signals. In *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, pages 380–382, 2021.

[37] Shinan Liu, Xiang Cheng, Hanchao Yang, Yuanchao Shu, Xiaoran Weng, Ping Guo, Kexiong Curtis Zeng, Gang Wang, and Yaling Yang. Stars can tell: a robust method to defend against {GPS} spoofing attacks using off-the-shelf chipset. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 3935–3952, 2021.

[38] Yang Liu, Sihai Li, Qiangwen Fu, Zhenbo Liu, and Qi Zhou. Analysis of kalman filter innovation-based gnss spoofing detection method for ins/gnss integrated navigation system. *IEEE Sensors Journal*, 19(13):5167–5178, 2019.

[39] Gianluca Falco Marco Rao. Code Tracking & Pseudoranges. *InsideGNSS*, 2012. https://insidegnss.com/wp-content/uploads/2018/01/IGM_janfeb12-Solutions.pdf.

[40] Daniel Moser, Vincent Lenders, and Srdjan Capkun. Digital radio signal cancellation attacks: an experimental evaluation. In *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks, WiSec 2019, Miami, Florida, USA, May 15-17, 2019*, pages 23–33. ACM, 2019.

[41] Maryam Motallebighomi, Harshad Sathaye, Mridula Singh, and Aanjhan Ranganathan. Location-independent GNSS Relay Attacks: A Lazy Attacker's Guide to Bypassing Navigation Message Authentication. In Ioana Boureanu, Steve Schneider, Bradley Reaves, and Nils Ole Tippenhauer, editors, *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec 2023, Guildford, United Kingdom, 29 May 2023 - 1 June 2023*, pages 365–376. ACM, 2023.

[42] Pavlo Mykytyn, Marcin Brzozowski, Zoya Dyka, and Peter Langendoerfer. GPS-spoofing attack detection mechanism for UAV swarms. In *2023 12th Mediterranean Conference on Embedded Computing (MECO)*, pages 1–8. IEEE, 2023.

[43] Sashank Narain, Aanjhan Ranganathan, and Guevara Noubir. Security of GPS/INS based on-road location tracking systems. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 587–601. IEEE, 2019.

[44] NASA. Broadcast ephemeris data. https://cddis.nasa.gov/Data_and_Derived_Products/GNSS/broadcast_ephemeris_data.html. Accessed 2025-01-11.

[45] National Buisness Aviation Association. GPS Spoofing: Should Operators Be Concerned? https://nbaa.org/news/business-aviation-insider/202403/gps-spoofing-should-operators-be-concerned/. Accessed 2024-07-31.

[46] Tyler Nighswander, Brent Ledvina, Jonathan Diamond, Robert Brumley, and David Brumley. GPS software attacks. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, 2012.

[47] Juhwan Noh, Yujin Kwon, Yunmok Son, Hocheol Shin, Dohyun Kim, Jaeyeong Choi, and Yongdae Kim. Tractor beam: Safe-hijacking of consumer drones with adaptive GPS spoofing. *ACM Transactions on Privacy and Security (TOPS)*, 22(2):1–26, 2019.

[48] Cillian O'Driscoll and Ignacio Fernández-Hernández. Mapping bit to symbol unpredictability in convolutionally encoded messages with checksums, with application to galileo OSNMA. In *Proceedings of the 33rd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2020)*, pages 3751–3765, 2020.

[49] Mark L Psiaki, Brady W O'Hanlon, Jahshan A Bhatti, Daniel P Shepard, and Todd E Humphreys. Civilian GPS spoofing detection based on dualreceiver correlation of military signals. In *Radionavigation Laboratory Conference Proceedings*, 2011.

[50] Katarina Rados, Marta Brkic, and Dinko Begusic. Recent advances on jamming and spoofing detection in GNSS. *Sensors*, 24(13):4210, 2024.

[51] Aanjhan Ranganathan, Hildur Ólafsdóttir, and Srdjan Capkun. SPREE: a spoofing resistant GPS receiver. In Yingying Chen, Marco Gruteser, Y. Charlie Hu, and Karthik Sundaresan, editors, *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking, MobiCom 2016, New York City, NY, USA, October 3-7, 2016*, pages 348–360. ACM, 2016.

[52] Harshad Sathaye, Gerald LaMountain, Pau Closas, and Aanjhan Ranganathan. Semperfi: Anti-spoofing GPS receiver for UAVs. In *Network and Distributed Systems Security (NDSS) Symposium 2022*, 2022.

[53] Harshad Sathaye, Maryam Motallebighomi, and Aanjhan Ranganathan. Galileo-SDR-SIM: An Open-Source Tool for Generating Galileo Satellite Signals. In *Proceedings of the 36th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2023)*, pages 3470–3480, 2023.

[54] Harshad Sathaye, Martin Strohmeier, Vincent Lenders, and Aanjhan Ranganathan. An experimental study of {GPS} spoofing and takeover attacks on {UAVs}. In *31st USENIX security symposium (USENIX security 22)*, pages 3503–3520, 2022.

[55] Desmond Schmidt, Kenneth Radke, Seyit Camtepe, Ernest Foo, and Michał Ren. A survey and analysis of

the GNSS spoofing threat and countermeasures. *ACM Computing Surveys (CSUR)*, 48(4):1–31, 2016.

[56] Daniel P Shepard, Jahshan A Bhatti, Todd E Humphreys, and Aaron A Fansler. Evaluation of smart grid and civilian UAV vulnerability to GPS spoofing attacks. In *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012)*, pages 3591–3605, 2012.

[57] Daniel P. Shepard, Todd E. Humphreys, and Aaron A. Fansler. Evaluation of the vulnerability of phasor measurement units to GPS spoofing attacks. *Int. J. Crit. Infrastructure Prot.*, 5(3-4):146–153, 2012.

[58] Fredrik Svanström, Cristofer Englund, and Fernando Alonso-Fernandez. Real-time drone detection and tracking with visible, thermal and acoustic sensors. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7265–7272. IEEE, 2021.

[59] Peter F Swaszek, Richard J Hartnett, Matthew V Kempe, and Gregory W Johnson. Analysis of a simple, multi-receiver GPS spoof detector. In *Proceedings of the 2013 international technical meeting of the institute of navigation*, pages 884–892, 2013.

[60] Ádám Szüllő, Rudolf Seller, Dániel Rohács, and Péter Renner. Multilateration-based UAV detection and localization. In *2017 18th International Radar Symposium (IRS)*, pages 1–10. IEEE, 2017.

[61] Çağatay Tanıl, Samer Khanafseh, Mathieu Joerger, and Boris Pervan. Kalman filter-based INS monitor to detect GNSS spoofers capable of tracking aircraft position. In *2016 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 1027–1034. IEEE, 2016.

[62] Nils Ole Tippenhauer, Christina Pöpper, Kasper Bonne Rasmussen, and Srdjan Capkun. On the requirements for successful GPS spoofing attacks. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 75–86, 2011.

[63] UT News. UT Austin Researchers Successfully Spoof an $80 million Yacht at Sea, 2013. https://news.utexas.edu/2013/07/29/ut-austin-researchers-successfully-spoof-an-80-million-yacht-at-sea/.

[64] Juan R. Vasquez. *Detection of spoofing, jamming, or failure of a global positioning system (GPS)*. PhD thesis, Air Force Institute of Technology, 1992.

[65] Kyle Wesson, Mark Rothlisberger, and Todd Humphreys. Practical cryptographic civil GPS signal authentication. *NAVIGATION: Journal of the Institute of Navigation*, 59(3):177–193, 2012.

[66] Kyle D Wesson, Daniel P Shepard, Jahshan A Bhatti, and Todd E Humphreys. An evaluation of the vestigial signal defense for civil GPS anti-spoofing. In *Radionavigation Laboratory Conference Proceedings*, 2011.

[67] Taiki Yamada, Katsutoshi Itoyama, Kenji Nishida, and Kazuhiro Nakadai. Sound source tracking by drones with microphone arrays. In *2020 IEEE/SICE International Symposium on System Integration (SII)*, pages 796–801. IEEE, 2020.

[68] Der-Yeuan Yu, Aanjhan Ranganathan, Thomas Locher, Srdjan Capkun, and David Basin. Short paper: Detection of GPS spoofing attacks in power grids. In *Proceedings of the 2014 ACM conference on Security and privacy in wireless & mobile networks*, pages 99–104, 2014.

[69] Kexiong Curtis Zeng, Shinan Liu, Yuanchao Shu, Dong Wang, Haoyu Li, Yanzhi Dou, Gang Wang, and Yaling Yang. All your GPS are belong to us: Towards stealthy manipulation of road navigation systems. In William Enck and Adrienne Porter Felt, editors, *27th USENIX Security Symposium, USENIX Security 2018, Baltimore, MD, USA, August 15-17, 2018*, pages 1527–1544. USENIX Association, 2018.

A Tracking and Synchronization Errors

Spoofing accuracy depends on the positioning of the drones carrying the transmitters, and their clock synchronization. We assume centimeter-level accuracy in drone placement using Real-Time Kinematic GPS [15, 23, 28], and synchronization in the order of ±50 ns using a GPSDOs [10]. We also assume that any systematic error can be easily corrected (e.g., through calibration) and that only Gaussian noise is left. We quantitatively evaluated the impact of these two problems in simulation (with an area of 10.24 km^2, attackers at 5000 m distance, 1000 measurements per location, 1 m standard deviation on the position of the transmitters, and 15 m standard deviation on the pseudo-ranges, corresponding to $\approx$ 50 ns on the synchronization). Results are shown in Figure 15.

B Additional Experimental and Simulation Data

Table 2 shows the results from our experiment with real hardware described in Section 4 and Figure 9 in tabular form for improved readability.

Figure 17 plots the cumulative error of CUSUM for the above data, showing that it does not always grow as the bias changes over time.

Figure 16 plots the trajectories measured by moving receivers on an error map like Figure 4, showing that they remain consistent.

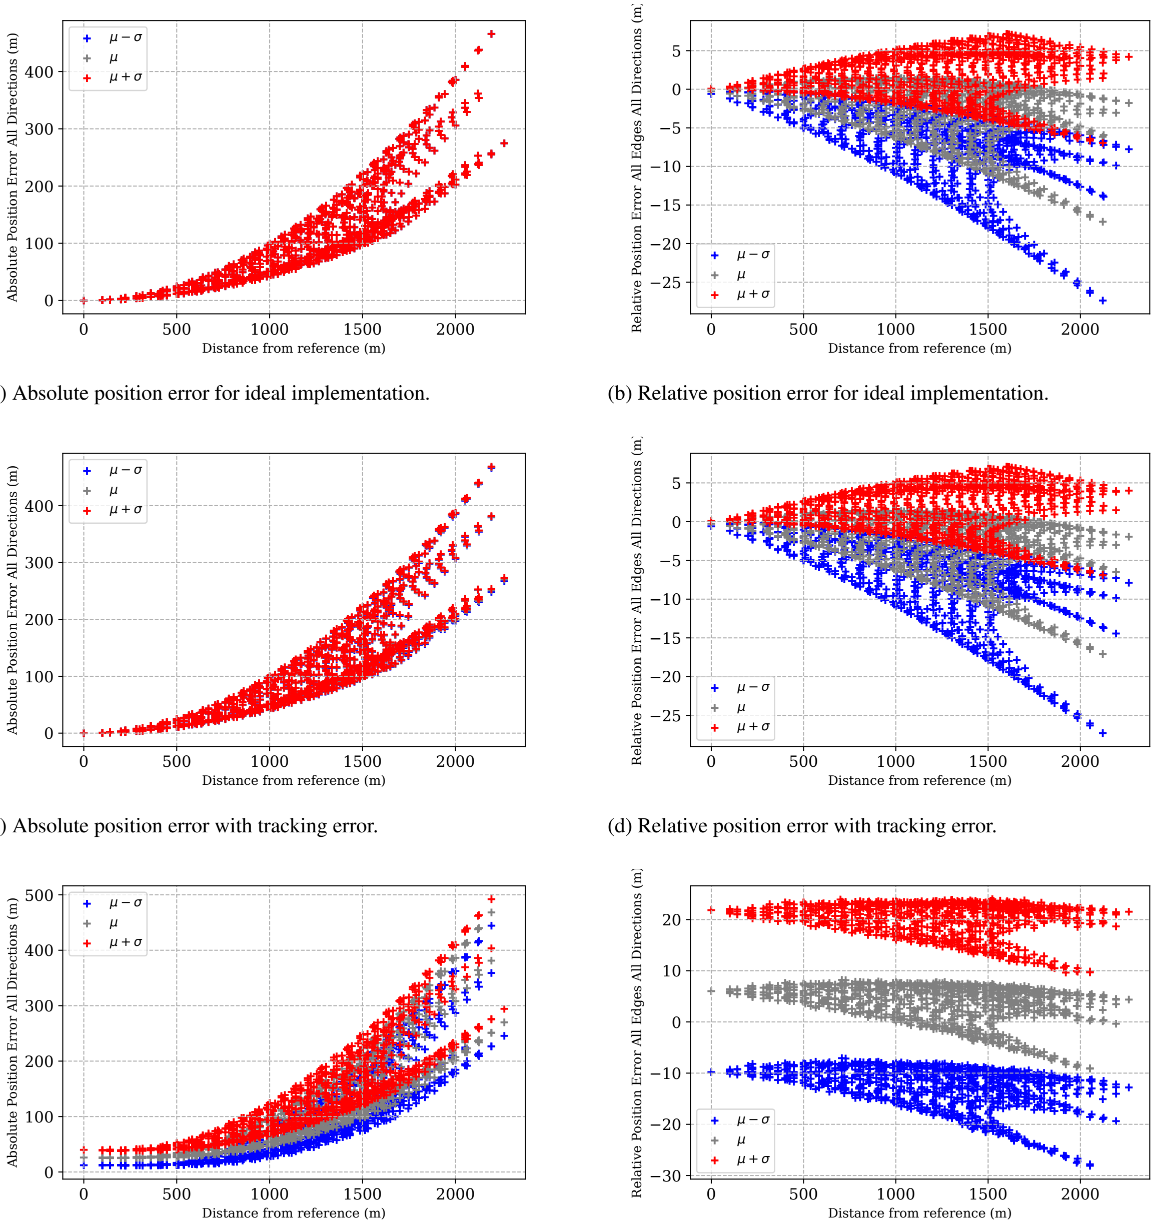

(a) Absolute position error for ideal implementation.

(b) Relative position error for ideal implementation.

(c) Absolute position error with tracking error.

(d) Relative position error with tracking error.

(e) Absolute position error with synchronization error.

(f) Relative position error with synchronization error.

Figure 15: **Synchronization and Position Errors.** Absolute (a,c,e) and relative (b,d,f) position error as a function of the distance from the reference. The absolute error is the magnitude of the difference between measured position and ground truth. The relative error is computed along 6 different directions of a 100 m by 100 m square formation, and represents the measured distance minus the ground truth. In an ideal implementation (a,b) there is no additional error apart from the bias introduced by `GNSS-WASP`. The spread in the curves is due to the bias being different along different directions. In the second case (c,d) we simulate an error of 1 m standard deviation in the position of the attackers tracking the satellites, without significant impact. In the third case (e,f) we simulate an error of 50ns standard deviation in the synchronization (i.e., 15 meter in the pseudo-ranges). For each case we show the mean μ and the interval $\mu \pm \sigma$ where σ is the standard deviation. The errors caused by synchronization are larger, but still relatively small compared to the 100 m edges of the receiver constellation, and the typical uncertainty in GPS measurements. Better GPSDO models could achieve lower jitter, and in our experimental results in Table 2 we generally have lower uncertainty than in this simulated model. At the same time, we also observed de-synchronization issues at high temperatures, which we reported in Figure 14 and discussed as limitation in Section 6.

		Ground Truth (m)	No Spoofing	Distance of Constellation from Reference (m)			
				10	**250**	**500**	**1000**
Receivers Pairs	**1 - 2**	11.02	11.23 ± 6.59	11.43 ± 5.66	11.19 ± 4.83	10.64 ± 4.32	11.67 ± 23.76
	1 - 3	17.02	17.21 ± 7.74	16.83 ± 6.12	16.46 ± 2.66	15.67 ± 2.99	15.71 ± 5.03
	1 - 4	14.56	14.49 ± 4.32	11.12 ± 8.70	12.94 ± 3.19	14.17 ± 1.33	17.93 ± 20.56
	2 - 3	11.62	11.44 ± 8.69	9.87 ± 6.06	11.58 ± 1.91	9.77 ± 5.43	13.01 ± 24.15
	2 - 4	21.53	21.32 ± 4.60	15.94 ± 17.52	21.55 ± 2.23	18.01 ± 5.53	25.63 ± 35.56
	3 - 4	18.52	18.24 ± 4.33	15.40 ± 16.11	19.57 ± 4.50	13.88 ± 2.01	19.26 ± 9.03

Table 2: **Distance between receiver pairs compared to ground truth in a no-spoofing scenario and under attack.** Measurements for no-spoofing scenarios are gathered from real receivers. Represented values are mean and confidence interval $(\mu \pm 3\sigma)$.

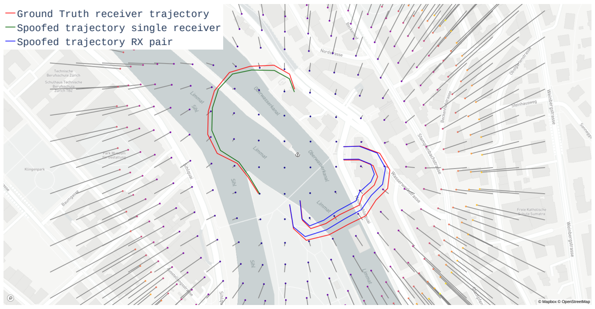
Ground Truth receiver trajectory
Spoofed trajectory single receiver
Spoofed trajectory RX pair
© Mapbox © OpenStreetMap

Figure 16: **Trajectories.** Under spoofing, trajectories closely match the ground truth. The attack bias does not significantly distort the acceleration/direction, like for relative distances. Hence, GPS data would be consistent with noisy INS data. Spoofing signal change naturally with the victim's position relative to the transmitters, without attacker's intervention.

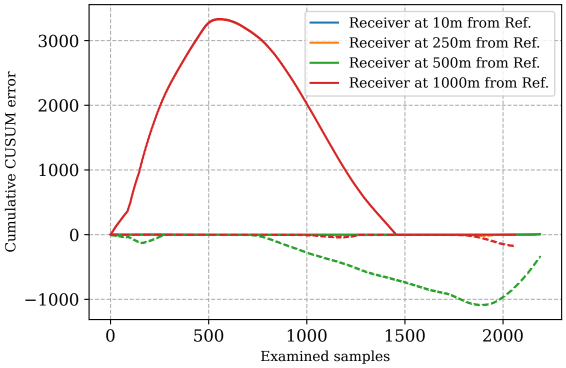

Figure 17: **CUSUM cumulative error over time.** Growth of the cumulative CUSUM error for the experiment with four receivers. FRR and FAR for different thresholds are given in Figure 13. Here we show that (i) it takes time before the error grows beyond a certain level increasing the latency of detection, and (ii) the error does not always diverge from zero (e.g., red line) reducing the detection capabilities of the method. This is because the bias induced by `GNSS-WASP` actually changes over time as satellites move.