An Experimental Study of GPS Spoofing and Takeover Attacks on UAVs

Harshad Sathaye*, Martin Strohmeier†, Vincent Lenders†, Aanjhan Ranganathan*
*Northeastern University, Boston, USA
† armasuisse Science + Technology, Thun, Switzerland

Abstract

Today, there is limited knowledge about the behavior of UAVs under GPS spoofing attacks in a real-world environment, in particular considering the interplay between the UAV's software as well as other equipped navigation aids and vision sensors. This work aims to understand the feasibility and requirements of fully controlling a UAV's movements by spoofing GPS signals alone. We enumerate the challenges in accomplishing a complete UAV takeover through GPS spoofing and controlling it without crashing. We design and implement a Real-time GPS Signal Generator (RtGSG) that can be configured to generate any arbitrary trajectory and is capable of making changes to GPS signals in real-time through user input, e.g., using a keyboard or joystick. We evaluate RtGSG on popular commercial UAVs from DJI and Autel through over-the-air spoofing experiments in a controlled chamber. We explore generic and UAV-specific GPS spoofing strategies in order to best achieve complete maneuvering control (e.g., velocity and direction). This work highlights that, although COTS UAVs remain vulnerable to GPS spoofing attacks, a complete takeover and control of the UAV requires careful manipulation of the spoofing signals in real-time. Finally, we release our implementation to the scientific community for further research.

1 Introduction

Today, there is a quickly increasing demand for unmanned aerial vehicles (UAV) across various civilian, military, and commercial applications, with market surveys [38] forecasting a doubling of the global retail UAV market in the next five years. Military and domestic law enforcement predominantly use UAVs for surveillance and reconnaissance operations. With their easy-to-use UAVs, manufacturers like DJI [20] and open-source platforms like ArduCopter [5] have enabled mass adoption of UAVs for civilian applications such as geographic surveys, photography, agriculture, recreational racing, package delivery, and many more.

This increased accessibility has also raised serious security and privacy concerns, especially after recent events in which civilian and military establishments were attacked using a slew of low-cost UAVs. For example, Heathrow and Gatwick airports reported several UAVs entering their airspace, significantly disrupting the air traffic for several days [54, 68]. There have been reports of terror groups using consumer UAVs laden with explosives to attack critical oil facilities and an airport in the Middle East [18, 49, 74]. Moreover, given these UAVs' low-visibility profile and cross-section, conventional air traffic radar systems are ineffective against these threat vectors. This has spawned a cottage industry of counter UAV systems, which promise reliable detection and protection against intrusions.

In general, including most of the above threat scenarios, UAVs heavily rely on the Global Positioning System (GPS) for positioning and navigation, particularly where they need to operate autonomously or in a pre-programmed fashion. GPS is an integral part of onboard decision-making that relies on positioning and navigation systems. Hence, GPS is seen as a single-point of failure for UAVs. At the same time, GPS has long been known to be vulnerable to jamming and spoofing attacks. The vulnerability can be profound: GPS spoofing provides an attack vector that enables control over the target UAV without compromising the flight control software or the command-and-control radio link. Furthermore, a GPS spoofing attack can be carried out by an attacker that is equipped with an RF transmitter. Since the attacker can generate spoofing signals for any arbitrary location, an attacker's proximity to the target is limited only by the attacker's amplification capabilities. An attacker equipped with a powerful enough transmitter or directional antennas need not be in close proximity of the target.

Widely-reported demonstrations show the feasibility of diverting unmanned ships [55], cars [59], and aerial vehicles [70]. In contrast to these attacks, GPS spoofing has also been explored as an active defense strategy, e.g., safely hijacking UAVs off a protected area [16, 57]. Despite the above demonstrations and the rapidly growing importance of UAVs,

there are limited studies on the feasibility of precisely controlling unmanned vehicles, specifically commercial off-the-shelf UAVs, by spoofing GPS signals. Prior work primarily focused on disrupting or altering the motion of the unmanned vehicle in a non-specific direction or performed the analysis on standalone GPS receivers. Kern et al. [40] used simulations to show the possibility of forcing the UAV in the desired direction by manipulating the GPS velocity in the opposite direction. However, this approach led to uncontrolled acceleration in the simulated environment. Noh et al. [57] provides a taxonomy of strategies to hijack consumer UAVs through GPS spoofing. However, the discussed hijacking approaches are limited to diverting the UAV in one direction, as they don't show the ability to maneuver the UAV, e.g., change the direction after the initial hijacking.

Importantly, no previous work has examined and field-tested such a controlled takeover of UAVs in a controlled real environment outside of a simulator. This state of affairs severely limits the available knowledge on the practicalities of GPS spoofing attacks on modern UAVs. GPS measurements are often fused with measurements from various sensors like inertial sensors, vision sensors, and distance measurement equipment. Given the tightly coupled nature of the system, it is vital to examine the UAV system as a whole.

Consequently, this work aims to understand the *feasibility* and the *requirements* of fully controlling a UAV's movements by spoofing GPS signals alone. We answer the following research questions:

1. Can an entity (adversarial or active defense) precisely control a UAV's movement by spoofing appropriate GPS signals?
2. What are the requirements and fundamental limitations of such spoofing strategies?

Specifically, we make the following contributions:

- We perform an exhaustive experimental analysis on the behavior of commercial-off-the-shelf (COTS) UAVs under a GPS spoofing attack. We execute our over-the-air spoofing experiments in a 15.24 x 15.24 x 6.7 m anechoic chamber equipped with a state-of-the-art motion capture (MoCap) system from OptiTrack [35] that offers precision tracking. Our setup enables us to characterize the response of the UAVs to different spoofing attacks for the first time in public literature. For experiments that require observing the UAV's behavior over longer distances, we use Arducopter [5].
- Based on our experiments, we enumerate several challenges in accomplishing a complete UAV takeover through GPS spoofing and controlling it without crashing. For example, even with the complete knowledge of the current state of the UAV, spoofing a pre-defined static location can cause the UAV to move in an unpredictable direction.
- We design and implement a Real-time GPS Signal Generator (RtGSG) that can be configured to generate any arbitrary trajectory and is capable of making changes to GPS signals in real time through user input. This enables us to modify the spoofing signal based on observing the UAV's reactions in real time, giving us better control of the UAV's trajectory and speed. RtGSG can interface with multiple software-defined radio frontends and can be controlled using any peripheral device like a joystick. Our signal generator can also interface with UAV simulators (like Arducopter) and UAV tracking systems (like OptiTrack), providing detailed analysis of the UAV motion. We will release our implementation for further research.
- We evaluate RtGSG on various UAVs from DJI and Autel and analyze the degree to which we can control the UAVs via GPS spoofing. We extract both generic and UAV-specific strategies to achieve complete maneuvering control. We were able to manually control and execute patterns like 180° turns through such a system as demonstrated in the video[1].
- Finally, we discuss limitations and highlight that COTS UAVs remain vulnerable (e.g., can be forced to crash or diverted away) to GPS spoofing. The complete takeover and control of the UAV is challenging and requires careful manipulation of the spoofing signals in real time.

The rest of the paper is organized as follows. In Section 2, we first describe the UAV ecosystem and provide a background on GPS and GPS spoofing attacks. In Section 3, we study the impact of conventional static-location spoofing and dynamic-path spoofing against consumer UAVs, present insights gained through these experiments, and lay down requirements for a complete takeover. This is followed by Section 4, where we implement and evaluate the real-time control strategies we develop based on the challenges and requirements identified. Then, in Section 5, we discuss the technical insights learned, the limitations, and the impact of our work. Finally, we provide an overview of the related work and conclude this paper.

2 Background

2.1 UAV ecosystem

UAVs are categorized as consumer, commercial, or military. With advancements in electronics and manufacturing, the lines between these categories are diminishing. For example, terror groups [61] have managed to make consumer UAVs combat-ready. Today, even COTS UAVs are capable of beyond visual line-of-sight operations with payload capacity from 500 g up

[1]Here is the link to a video demonstration of this attack. https://www.youtube.com/watch?v=EtaQ_BQFn-M

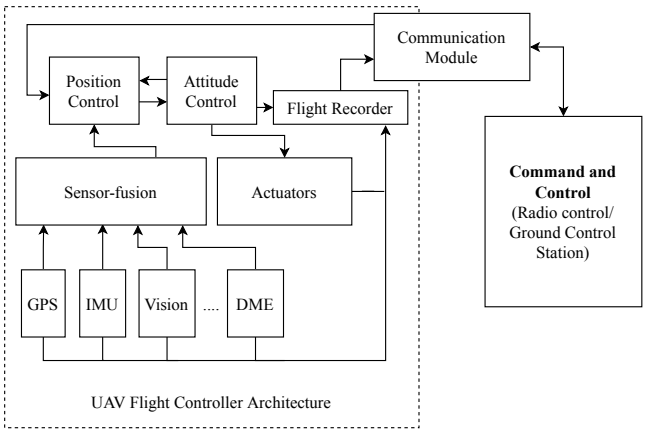

Figure 1: Schematic of a generic UAV flight controller architecture depicting various modules like sensor-fusion, position and attitude control, flight recorder, communications unit, and a battery of sensors.

to 200 kg [4], flight speed up to 70 kmph, and flight height of more than 5 km above see level. Irrespective of their application, UAVs generally implement the following architecture. The main components of a UAV system are the vehicle itself, the operator, a wireless radio controller, and a ground control station built specifically for managing autonomous flight. Powerful onboard microprocessors act as flight controllers capable of sensor fusion, navigation, advanced mission planning, and safety-critical decision-making. Refer to Figure 1 for a schematic representation of a generic flight controller and its various components. Autonomous flight requires a programmed mission that includes a pre-defined trajectory with waypoints where each waypoint of flight segment can have its speed and altitude profile. Even consumer UAVs come with a battery of sensors like GPS, vision sensors, inertial sensors (IMUs), and various types of distance measurement equipment (DME) that aid in navigation and position control to provide safe and efficient flight.

Typical UAVs implement a proportional, integral, and derivative (PID) controller for attitude and position control. It is ultimately responsible for driving the motors that generate thrust that moves the UAV as required. The flight controller uses the outputs of the PID controller to determine how fast the motors should spin to achieve and maintain the desired attitude and position. The sensor-fusion algorithm which can fuse measurements from various onboard sensors like GPS, inertial measurement units, and opti-flow sensors provides the input to this control loop. Some of the most widely adopted sensor-fusion algorithms are based on an extended Kalman filter (EKF).

Figure 2 shows a schematic of a typical PID controller implementation for horizontal position control. A typical PID

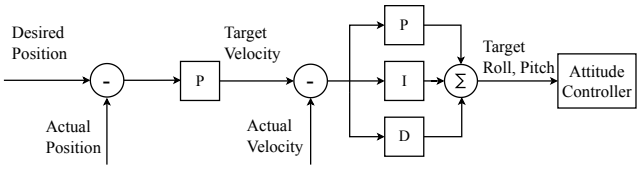

Figure 2: Schematic of a typical PID control sequence for position control. First, the desired and actual positions error is used to compute target velocity. Next, the errors calculated by subtracting target and actual velocity are used by the controller to set target roll and pitch. These target values govern how fast the motors should spin to achieve the desired position and attitude.

controller takes the form of:

$$u(t) = k_P e(t) + k_I \int_0^t e(\tau)\, d\tau + k_D \frac{de(t)}{dt} \tag{1}$$

where $u(t)$ is required change, $e(t)$ is the error in desired and actual values and k_P, k_I, and k_D are the respective gains. A UAV that is set to hover; i.e., the desired location is a fixed value, can experience a drift because of two factors: i) internal measurement errors that arise because of inertial sensors that drift and other faulty measurements; or ii) external factors like wind or someone picking up the UAV and moving it to a different position. When the UAV experiences such drifts, the PID controller issues appropriate commands to actuators that control the motors. This enables the UAV to maintain its position.

These features collectively enable UAVs to carry out fully autonomous flights. Moreover, vision sensors and distance measurement equipment also provide automatic obstacle detection and avoidance capability. Typically a UAV supports the following flight modes: i) Manual: The operator is in complete control of the vehicle, the flight controller does not provide any stability control; ii) Stabilize/Loiter: in this mode, the operator is responsible for position control and the flight controller assists in stabilizing the UAV by taking over when the operator does not provide any input and maintain altitude. Additionally, this mode usually has roll and pitch restrictions to maintain thrust. iii) Mission: A complete autonomous operations mode where the flight controller governs the attitude and the three-dimensional position; In this mode, the flight controller executes a predetermined mission, usually a set of waypoints in the form of GPS coordinates, with each flight segment having its speed and altitude profile; and iv) Land: This mode is usually activated at the end of a mission or as a failsafe mechanism. The UAV automatically lands at the current position or a pre-configured location in this mode, usually its takeoff or home location.

The flight controller has certain predetermined operations, called failsafes, that are triggered to ensure the vehicle's safety in case it encounters any errors in flight. These errors can

result from faulty sensor measurements, loss of thrust, a malfunctioning battery, or even high-speed winds. For example, Autel UAVs will abort the current mission and land when the battery runs out of charge [9]. Often such failsafes are user-configurable. However, there are some *terminal failsafes* that cannot be overridden, even through human intervention e.g., EKF variance in ArduCopter and no-fly zone (NFZ) restrictions in DJI drones.

2.2 GPS background

The Global Positioning System (GPS), the most widely used navigation system, consists of 29[2] operational satellites roughly at an altitude of 20,200 km. These satellites continuously transmit individual satellite ephemeris and timing data at 50 bps, allowing a receiver to localize itself with respect to known satellite positions. GPS provides a civilian positioning service with accuracy up to 5 m on the L1 frequency [77]. In this service, each satellite is assigned a unique publicly available coarse-acquisition (C/A) code that enables the receiver to track and decode signals from different satellites transmitting on the same carrier frequency through code-division multiple access technology. The GPS receiver apparatus includes an antenna and a signal processing chip that outputs a variety of information, including; the receiver's positionand the estimated altitude and velocity of the receiver $\boldsymbol{v}$ decomposed into Easting v_e and Northing v_n velocity in m/s. While UAVs support multiple satellite navigation constellations, GPS is the typical constellation used across all UAV platforms and manufacturers.

2.3 GPS spoofing attacks

Due to the lack of any form of authentication and public access to satellite spreading codes, modulation techniques, and data structure, GPS is vulnerable to signal spoofing attacks which are physical-layer attacks where the attacker transmits a pre-crafted signal that contains appropriate satellite messages. When the receiver uses these counterfeit signals, the receiver calculates the position, navigation, and timing (PNT) solution initially programmed by the attacker. This deceives the receiver into believing that it is at the location spoofed by the attacker rather than its actual position. An attacker can achieve this by either manipulating the navigation messages or modifying the time of arrival of these messages. Additionally, an attacker can reuse current navigation messages to make the attack stealthier.

Broadly, there are two ways of hijacking a target GPS receiver. In the first method, an overshadow attack, the attacker transmits fake GPS signals with enough power to bury the legitimate signals under the noise floor. A receiver can easily detect such an attack because of a sudden loss of lock. The second way is a more stealthy approach. The attacker first synchronizes with the legitimate satellite signals. Once it is synchronized, it increases the power of its signal and then slowly starts adding code offsets that move the receiver away from its actual location. In [72], the authors provide requirements for executing such an attack. In both these attacks, the attacker's objective is to hijack the receiver and deceive it into believing it is at a location of the attacker's choosing.

[2] As of January 1, 2022 - [2]

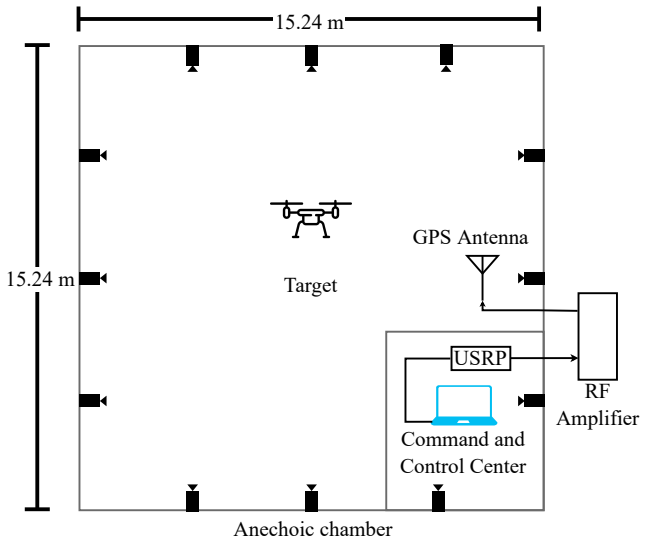

Figure 3: A layout of the anechoic chamber equipped with motion capture system, GPS signal generator, and the transmission antenna. The RF amplifier is installed outside the chamber to minimize EM interference inside the chamber.

2.4 Attacker Goals and Assumptions

In our work, we consider an attacker capable of generating and transmitting GPS signals. In attacking a UAV, an attacker's main goal is to force the UAV to move to a specific location by spoofing GPS signals The UAV is assumed to be within the attacker's radio range and is able to receive the spoofing signals. We also assume that the attacker has managed to takeover the UAV's GPS receiver by either a seamless takeover attack, as explained in [63, 72], or through a non-coherent overshadow attack. Prior work has extensively analyzed the spoofing vulnerability of standalone GPS receivers [39, 56, 62, 72]. The received signal strength of the GPS signals on ground is typically around -127.5 dBm and, hence, it is trivial for an attacker to overshadow the legitimate signal with the adversarial signal.

Researchers have also demonstrated the ability to steer yachts [55], cars [59], and drones [57, 70] to some extent through various GPS spoofing experiments. In this work, we focus on evaluating the UAV's response to GPS spoofing and strategies to fully control the *UAV*. In the following sections, we describe the limitations of static location and dynamic path

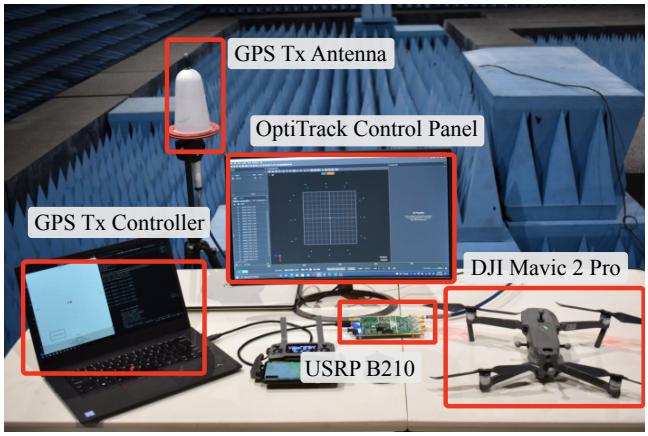

Figure 4: A photo of our actual setup featuring RtGSG, OptiTrack control panel, GPS Tx antenna, and DJI Mavic 2 Pro, one of the target UAVs.

Figure 5: All UAVs that we used in our study. From top left, i) Autel EVO II, ii) DJI Mavic Mini, iii) DJI Mavic Pro, iv) DJI Mavic Air 2, and v) DJI Mavic 2 Pro.

spoofing attacks that rely on pre-crafted signals on UAVs and explore the possibility of asserting fine-grained control over the UAV based on UAV's retroactions to GPS spoofing and the effect of these retroactions on the process of a complete takeover.

3 Evaluation of Conventional GPS Spoofing Attacks

The goal of this section is to categorize and analyze the response of UAVs to pre-defined static-location spoofing and dynamic-path spoofing. Specifically, we analyze the challenges and limitations of GPS spoofing and specify requirements to gain complete control of the target UAV. It is important to note that in this work we focus on evaluating the UAV's response to GPS spoofing and strategies to takeover the *UAV* and not the receiver.

3.1 Evaluation Setup

Transmitting GPS signals over the air in an uncontrolled setting is illegal. We perform over-the-air GPS spoofing experiments in a 15.24 x 15.24 m shielded anechoic chamber that provides more than 100 dB of attenuation. Given the tight bounds of the chamber, we are limited to spoofing experiments over shorter distances. The building materials used in the chamber construction is a source of strong magnetic interference, and hence the UAV requires constant calibration. Despite the trade-off between safety and realism, the shielding enables us to transmit GPS signals without running into legal issues and without the need for tethered UAV operations. Moreover, environmental factors that can affect UAV's performance, like wind and temperature, are virtually non-existent inside the chamber. Reducing the effect of environmental factors ensures that the UAV's motion is affected only by the spoofed GPS locations, creating a best case setting for evaluating attacker's requirements to execute a UAV takeover through GPS spoofing.

The anechoic chamber is equipped with a motion capture system that runs 24 OptiTrack cameras that can track objects with mm precision and are capable of providing live tracking data at 120 Hz [35]. It is important to note that the motion capture system is only used for tracking and recording the UAV's motion. A GPS signal generator [25] with a USRP B210 as the RF frontend was connected to a Ophir 5293 RF amplifier [58] that supports output power upto 50 W. The output of the RF amplifier is fed to a ETS-Lindgren's Model 3181 [27] omnidirectional antenna. Refer to Figures 3 and 4 for the schematic and the actual photo of our test setup.

We evaluate our attacks on UAVs manufactured by DJI [20] and Autel [8], shown in Figure 5. DJI and Autel are two leading consumer and commercial UAV manufacturers, with almost 76% market share owned by DJI alone [36]. For tests where the primary metric is distance, we used popular COTS UAV simulator software that runs ArduCopter [5] firmware alongwith Gazebo [64], an advanced physics and environment simulator.

3.2 Preliminary Observations

Fallback sensors and non-GPS navigation: Modern UAVs are equipped with vision sensors that can provide positioning information accurate up to 0.3 m horizontally and 0.1 m vertically [22]. UAVs typically fallback to vision positioning system in a GPS-denied environment. We began our experiment by placing the UAV in the center of the test area and instructed the UAV to take-off and maintain an altitude of 2 m. Once we visually verified that the UAV was stable we started introducing motion to the generated signals. In

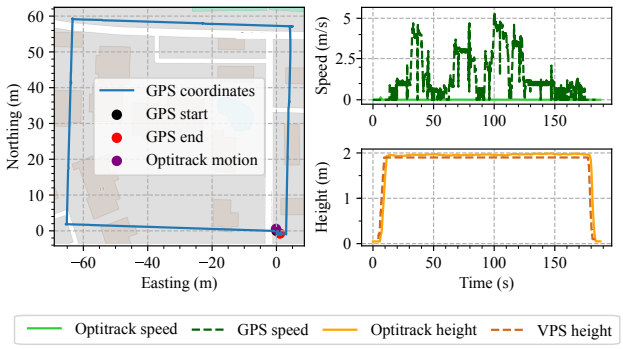

Figure 6: Comparison of UAV's GPS measurements and its actual motion. In spite of introducing a motion that takes the GPS receiver for a ride, the UAV manages to hover. Ground speed as calculated by the UAV's GPS receiver shows a maximum speed of 5.4 m/s while the ground speed calculated from OptiTrack data is constant at 0 m/s.

this experiment, the spoofed GPS signal introduces a motion such that the receiver believes it is moving along a path 254 m long with a maximum speed of 5.4 m/s. In spite of introducing this type of motion, the target UAV did not budge and hovered steadily in-place. It was only after we turned off the vision sensors that the UAV reacted violently with rapid acceleration to our spoofing. Figure 6 shows the result of one such test where one can clearly see the UAV's position (as tracked by the motion capture system) being stable as opposed to a change in GPS measurements. From this experiment, we conclude that the target UAV was in fact prioritizing vision sensor measurements for positioning and navigation over GPS measurements. This shows that a UAV can survive a GPS spoofing attack by relying on other available sensors.

However, there are some limitations associated with vision sensors; these sensors require optimal lighting conditions and can provide accurate guidance only up to an altitude of a few meters. For example, DJI Mavic 3, the latest UAV from DJI, provides vision positioning only up to a height of 18 m and with flight speed < 6 m/s [23]. To evaluate the effect of GPS spoofing on a UAV switching from vision to GPS positioning, we disabled the downward vision sensors in-flight to simulate a scenario where the UAV is flying at an altitude greater than 18 m. As soon as we disabled the downward vision sensors, the UAV reacted by accelerating rapidly, eventually crashing into the RF energy-absorbing foam. This incident prompted us to find a suitable target velocity that will allow us to observe the UAV's reaction without creating a safety hazard. We tested multiple target velocities and narrowed down to a suitable velocity using the UAV simulator. For this we used a trial and error method for different velocity configurations. The initial acceleration of the UAV is directly related to the spoofed GPS velocity and is evident from Figure 7.

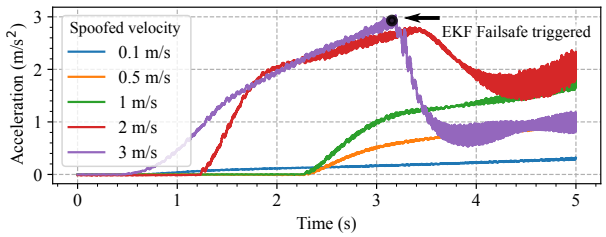

Figure 7: Results of first 5 s of the UAV's reaction to different spoofed GPS velocities. For higher GPS velocities, the UAV's more aggressive response is evident from the steep rise in the acceleration values.

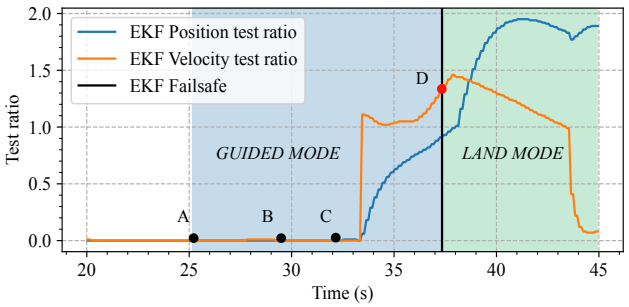

Figure 8: EKF variance test ratios as calculated by ArduCopter. At point A, the UAV arms and starts to takeoff, at point B, it completes takeoff. At point C, the attacker starts spoofing west at 2.5 m/s and at point D, the UAV activates a terminal failsafe and permanently switches to *LAND* mode.

Terminal failsafes: Some autopilot software like ArduCopter and PX4 implement what we define as a *terminal failsafe*. When such a failsafe is activated, the UAV switches to *LAND* mode. The flight controller calculates the position and velocity test ratios using EKF innovations after sensor fusion and triggers a failsafe if these test ratios exceed a predetermined threshold [6]. Figure 8 shows the effect of GPS spoofing on EKF test ratios where GPS velocity is set to 2.5 m/s. In ArduCopter, when the flight controller switches to *LAND* mode, it still tries to maintain horizontal position by relying on GPS. Depending on the altitude of the UAV, the attacker has very limited time to further control it.

3.3 Impact of Spoofing a Static Location

Despite of all the sensors and the non-GPS navigation systems that can be incorporated in a UAV, GPS still remains the most important navigation system. And unlike non-GPS systems, it also poses a greater threat. Even a naive adversary that can only transmit static location can still cause considerable damage to the UAV. In this experiment, we evaluate the final bearing of the UAV with respect to its take-off position and

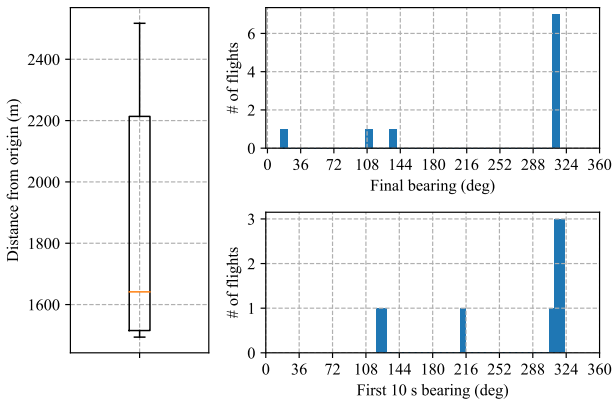

Figure 9: A comparison of 10 simulated flights where the attacker executes a naive static single location spoofing attack. The box plot shows the statistics of the distance covered by each flight before a terminal failsafe is activated.

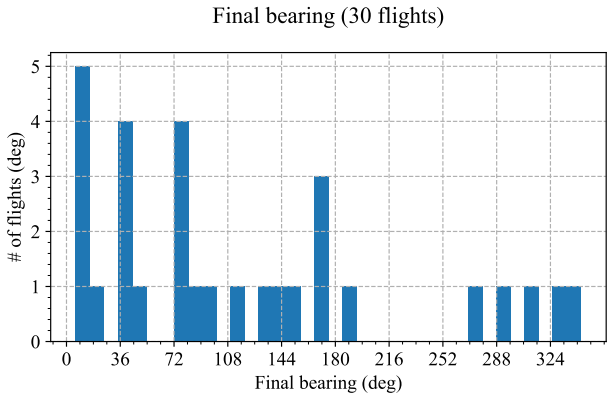

Figure 10: A comparison of 30 flights with a naive attacker. The target of these attacks is a DJI Mavic 2 Pro. The final bearing of these flights show that the UAV's behavior is unpredictable and uncontrollable.

distance that the UAV travels before it looses thrust or till any failsafes are activated. These experiments are conducted in a simulator and in real-world settings.

In this experiment, the UAV is programmed to hover at a certain location. The attacker spoofs the UAV's actual location, a single static location, i.e., the spoofed location remains unchanged throughout the attack. The objective of the attacker in this type of an attack can be to force the UAV to stop and hover. Even though this seems benign, in our experiments we found that the UAV's response is unpredictable and uncontrollable. Flight controllers use EKF-based sensor fusion algorithms for state estimation, which provides the UAV with increased stability during flights. The UAV's uncontrolled movement can be attributed to the lack of the correction that is required to control the drift and biases that develop in inertial measurements. Because of the IMU drifts, the position and velocity estimates obtained from EKF differ from GPS measurements. As a direct result of the discrepancy in these two measurements, the flight controller accelerates to compensate for the difference. Since the PID controller implements a feedback loop, the errors propagate and force the flight controller to make more drastic corrections.

To execute this attack, we configured the GPS signal generator to transmit a fixed static location. ArduCopter starts drifting and eventually triggers an EKF variance failsafe as result of position and velocity error accumulation. Once the flight controllers trigger the EKF failsafe, the UAV switches to *LAND* mode and aborts any ongoing mission. For this experiment, our evaluation metrics are the distance the UAV travels before an EKF failsafe is triggered, the final bearing of the UAV, and the bearing in the first 10 seconds of the flight. The difference in the final bearing and in the first 10 s shows how unpredictable such an attack can be. The results of these flights are summarized in Figure 9. The average flight distance was 1861.83 m with a standard deviation of 406.39 m.

This shows that not only is the direction random and uncontrollable, but the distance it covers is also unpredictable. This makes such an attack very unreliable, especially if the attacker requires the UAV to reach a specific location. A similar experiment was performed on DJI Mavic 2 Pro; given the space constraints of the anechoic chamber, the evaluation metric was just the bearing. The results of this experiment are summarized in Figure 10. Unlike the results of the simulator, the real UAV shows a lot of variation in terms of final bearing. Real sensors are deeply affected by environmental factors and often require re-calibration for normal operations. Such factors generally do not apply to the built-in configurations available in simulators, illustrating the limits of their utility.

3.4 Impact of Spoofing a Dynamic Path

In this experiment, our goal is to analyze and understand the behavior of a UAV subject to dynamic-path spoofing. We evaluated this attack entirely on real UAVs with live over-the-air GPS signals. After take-off the UAV is set to hover at its current location. The objective of the attacker is to transmit a signal that forces the UAV to move away from its current position in a direction of attacker's choosing. In this attack scenario, we pushed the GPS receiver away from its original position by generating a spoofing signal that moves in a specific direction.

In a dynamic-path spoofing scenario, after a successful GPS takeover, the attacker adds velocity to the spoofed locations and deceives the UAV into perceiving that it is moving with a heading of $\alpha°$. This activates the attitude and position control mechanism and forces the UAV to move in the opposite direction i.e., $(\alpha - 180)°$. The UAV's reaction to such an attack is shown in Figure 11. Consider a UAV that is at

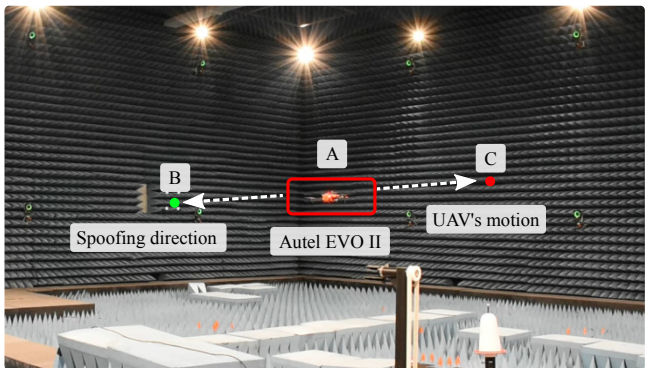

Figure 11: Target UAV's response to dynamic path spoofing. When subjected to a GPS signal that simulates a motion in the direction of point B, the UAV responds by moving in the direction of point C.

point A, the attacker introduces a GPS signal that deceives the UAV's GPS receiver into believing that it is moving towards point B. As a result, the UAV starts moving towards point C. For evaluating dynamic path spoofing scenario, we executed 5 flights in each of the 4 directions, i.e., north, south, east, and west w.r.t to the origin. Based on the results of our vision to GPS positioning transition experiment described in Section 3.2 and the tight space constraints, the magnitude of the spoofed velocity was set to 0.1 m/s. The sequence of each flight was as follows: i) the UAV takes off, ii) once stable, the operator switches to "GPS Only" mode to simulate higher altitude, iii) the spoofer is activated, and iv) as the UAV gets closer to the walls, the operator intervenes and lands the UAV manually. Figure 12 shows the response of the target to a spoofed velocity vector $\boldsymbol{v}_{en} = [-0.1, 0]$ that forces the UAV to fly east. Figure 13 shows the error in final bearing of all 20 flights. In all these experiments, we observed that the UAV reacts as expected and goes in the expected direction with an average error of 2.56°. As a next step, we introduced a second change in the direction. Specifically, we changed the direction of the spoofed trajectory by 90°. This strategy showed limited success, only 3 out of 17 flights followed the required change in bearing. Moreover, without any velocity control the target flies a curve, making it impossible to achieve sharp turns because of the momentum that the UAV already developed due to the spoofing attack.

3.5 Key Insights and Lessons Learned

The UAV's response to GPS spoofing can be attributed to the correction maneuver enforced by the position and attitude control described in Section 2.1. The PID controller responds to the changes in the UAV's actual position and velocity measurements derived from sensor fusion by providing control inputs to compensate for the error between the desired and the actual position. Recall Equation (1), over a period of time as

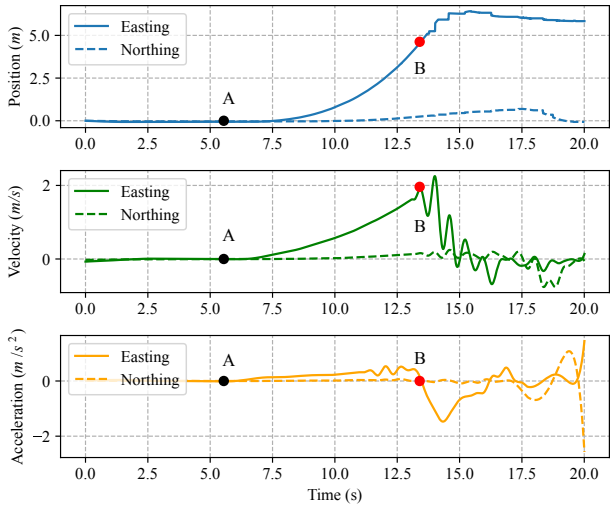

Figure 12: Position, velocity and acceleration data for the takeover of a DJI Mavic 2 Pro. The position data was obtained from a motion capture system. At point A we start introducing a motion as $\boldsymbol{v}_{en}$ and at point B the operator intervenes once the UAV gets closer to the wall.

a result of the integral ($k_I \int_0^t e(\tau)\, d\tau$), the errors are magnified and the corrections to even small errors get more aggressive. As a result, the UAV tries harder to overcome the error. Since the spoofed location is consistently going away from the original position, the UAV keeps increasing its velocity as a result of aggressive corrections explained earlier. In Figure 12, the target's velocity keeps increasing until the operator takes over.

The UAV achieves the required acceleration by manipulating its pitch and roll. The thrust value $\boldsymbol{T}$ required to maintain its altitude and to stay afloat is given by

$$T = \frac{mg}{cos\theta cos\phi} \tag{2}$$

where θ and ϕ are pitch and roll, respectively. According to its dynamics, the quadcopter fulfills the acceleration requirement by manipulating the pitch and the roll. As the magnitude of corrections increases, the magnitude of required acceleration also increases. In order to achieve the desired acceleration, the UAV tilts[3] so much that it is no longer able to generate enough thrust to keep itself afloat. As a result, it will lose altitude and crash. From this, we conclude that even if the UAV goes in the specified direction with minimal directional errors, the attacker's control over the target is limited only to the direction in which the UAV goes, and the attacker has no control over UAV's speed. From these experiments we understand that:

- Spoofed GPS velocity induces acceleration in the UAV

[3] A combination of pitch and roll.

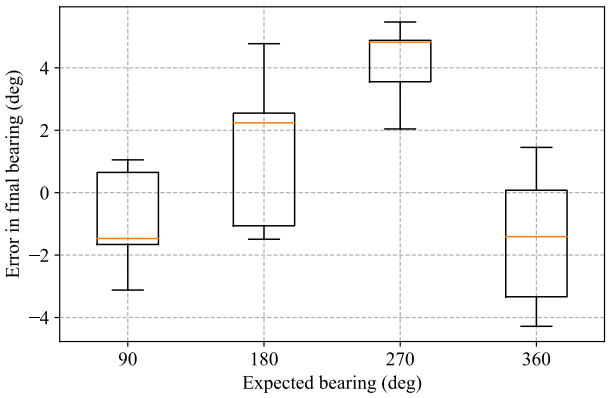

Figure 13: A boxplot of error in final bearing against the expected bearing. Five flights were carried out for each expected bearing. A trajectory was generated as described in Section 3.

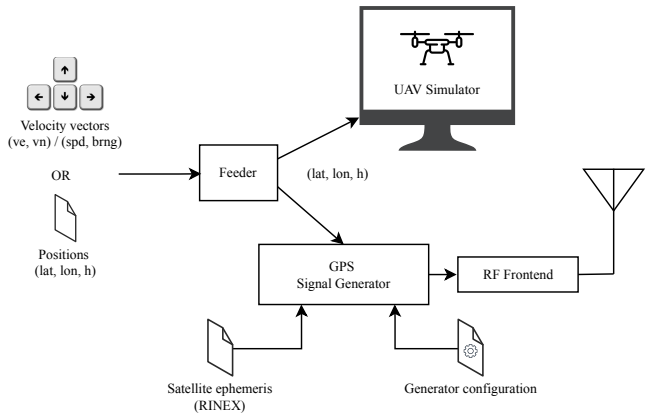

Figure 14: A schematic of RtGSG's architecture showing the feeder capable of using positions or velocity vectors with an optional link to send the coordinates directly to a UAV simulator, the GPS signal generator that generates raw IQ samples and interfaces with an RF frontend capable of transmitting the generated GPS signals.

- A UAV will continue accelerating in the initially spoofed direction until it runs out of battery or crashes
- Complete control of the UAV requires direction control as well as speed control

As explained Section 3.4, in such an attack the attacker can *only* control the *direction* and not the speed. Even if the attacker changes the direction of spoofed trajectory, the attacker is unable to change the direction of the UAV.

Consequently, we establish that the attacker should be able to force the UAV to execute four specific maneuvers to constitute a complete takeover of the UAV. These are i) flight with constant direction, ii) flight with constant velocity, iii) flight with variable direction, e.g., the ability to make tight turns in an environment with strict mobility constraints like an urban setting, and iv) land. These specific maneuvers ensure complete control over the direction of the UAV and the distance it travels. The attack strategies proposed by various researchers in the past do not fulfill these requirements, making the proposed attacks uncontrollable. Moreover, terminal failsafes like EKF variance bounds make it even more difficult to effectively exert control over the target UAV. We identified that for complete takeover we need: i) a GPS signal generator capable of user-controlled real-time trajectory manipulation and ii) a strategy for controlling the acceleration of the UAV. In the following section, we outline the strategies that we developed to address the challenges mentioned above through experiments.

4 Real-time Control of UAV via GPS Spoofing

Based on the insights described above, we conclude the requirement for dynamically manipulating GPS signals in real time based on a UAV's response to the spoofing signal. In this section, we present our real-time GPS signal generator that addresses this need. We also propose and evaluate strategies for post-takeover direction and velocity control.

4.1 Real-time GPS Signal Generator

Commercially available off-the-shelf hardware and open-source software GPS signal generators are often limited to predetermined trajectories, i.e., they are not capable of user-controlled real-time trajectory manipulation. This is a requirement that we identified in the previous section. To facilitate this requirement, we built the *Real-time GPS Signal Generator* (RtGSG), a GPS signal generator system that allows real-time trajectory manipulation. This system is based on an open-source GPS signal simulator, *GPS-SDR-SIM* [25]. RtGSG comprise three main components as shown in Figure 14: i) feeder, ii) GPS signal generator, and iii) RF front-end.

The feeder can accept a set of predefined trajectories in the form of a time series that contains positions $P = \{p_1, p_2, p_3, .., p_n\}$ where $p_i = (latitude, longitude, height)$ or an initial location and velocity vector[4] $\boldsymbol{v}_{en}$ as $[ve, vn]$ (*Easting and Northing*) or as speed and bearing. The feeder can also accept velocity vectors through human interface devices like a keyboard or a joystick. This is especially useful in our *Human-in-the-Loop* (HITL) spoofing system where a user can manipulate the GPS signal using a keyboard or joystick, similar to playing a video game. The feeder is responsible for computing the correct location and time and updating the GPS signal generator. The GPS signal generator receives the loca-

[4]We use the following units: all positions are in decimal degrees, velocity is in m/s and bearing is in degrees w.r.t to north, unless stated otherwise.

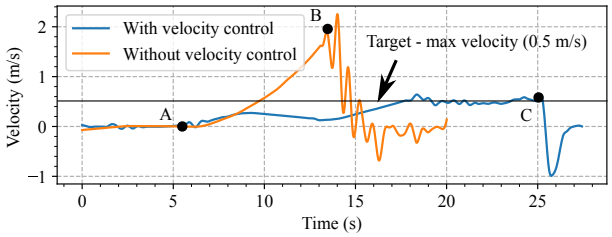

Figure 15: A comparison between two test flights with and without velocity control. At points B and C, the operator takes over. Without velocity control, the UAV keeps accelerating.

tions, constructs the navigation message using the supplied satellite ephemeris data in RINEX [31] format, modulates the message, and generates raw IQ samples. Following the generation of IQ samples, the RF front-end module interfaces with software-defined radios like USRPs [37] and LimeSDRs [46] that can transmit the generated samples in real time. It is important to avoid any type of hardware-dependent sample drops as they can cause the target receiver to lose the GPS lock. Sample generation and consumption should be synchronized to avoid buffering IQ samples that can lead to position/time jumps. These modules enable RtGSG to manipulate trajectories and transmit the generated signals on-demand, instantly, and precisely. We will release RtGSG to the scientific community for further research.

4.2 Strategy for Velocity Control

As demonstrated in Section 3.4, a conventional fire-and-forget attack with predetermined trajectories ensures that the UAV goes in a specific direction, but the lack of speed control makes the UAV uncontrollable. At a high level, our system achieves velocity control by deceiving the target into believing that the correction has worked. Then, due to the correction maneuver it executes, it approaches the "target position". However, since the integral term responds to errors from the past, it can often overshoot the target. A PID controller is designed to compensate for this and handle overshoots.

From the lessons learned in Section 3.5, GPS velocity induces acceleration in the target UAV, and the direction of the acceleration vector depends on the direction of the GPS velocity vector with respect to its original position. Hence, as the spoofed location approaches the target position, the UAV gradually decelerates. However, the UAV does not immediately stop its motion when the onboard receiver indicates it has arrived at the target position. As a result of initial spoofing, the UAV has already gained momentum and needs to overcome that. On the other hand, even if the attacker "jumps" to the original location, the UAV faces a similar issue, and hence the UAV continues its motion. Thus, the attacker needs to simulate a motion that resembles a UAV's correction maneuver.

In reality, the UAV can be considered as a black box, and hence the attacker does not have access to the correction model implemented by the target UAV. To overcome this issue, we came up with a strategy to identify the reaction time of the UAV. The attacker can learn the reaction time by observing the target's response during the course of the attack, specifically the time it takes to achieve the spoofed velocity. Alternatively, the attacker can also estimate this value using quantities such as maximum angular velocity, the maximum tilt angle, and the total thrust that the UAV can generate. Product specifications and data sheets make these values readily available for consumer and commercial UAVs. Moreover, consumer and commercial UAVs are mass-produced, and hence they have set standards that make variations amongst these values within a specific model unlikely. However, certain environmental factors and biases unique to individual sensor units may affect these values.

For a spoofed velocity of $\boldsymbol{v}_{at}$, the target UAV responds by tilting and accelerating to catch up with the GPS velocity. The UAV experiences a lag in achieving the target velocity. We define this lag as the UAV's reaction time. However, the UAV's objective is to correct the position error. Therefore, even if the UAV achieves the target velocity, it is still away from the target position, and hence it continues the correction. This reaction time provides us with the average acceleration of the UAV. We use this value to time the spoofed GPS signal's *return to launch* (RTL) maneuver where the spoofed trajectory starts moving back to the initial position rather than away from it. This is done by changing the direction of $\boldsymbol{v}_{at}$ by 180°, i.e., the new spoofed velocity $\boldsymbol{vl}_{at} = -\boldsymbol{v}_{at}$. This maneuver causes the UAV to decelerate. We identified that a well-timed spoofed GPS velocity-induced acceleration/deceleration routine could be used to control the velocity of the UAV — the key is to maintain an average acceleration of 0 m/s^2. In Figure 15, we show post-takeover velocity control in action by maintaining the UAV's velocity under a maximum target velocity of 0.5 m/s. To further analyze the effectiveness of this technique, we perform 48 flights that make use of the same reaction time value to control the velocity, and the results are shown in Figure 16. For this experiment, the post-takeover UAV velocity was set to 0.5 m/s. In 56.25 % of all 48 flights, we managed to maintain average acceleration below 0.015 m/s^2. For 81.25% of the flights, we were able to maintain an average $|\boldsymbol{v}|$ below the set maximum target velocity of 0.5 m/s with average observation time[5] of 19.4 s.

4.3 Strategy for Direction Control

As mentioned in Section 3.4, just changing the direction of the spoofed trajectory is not sufficient because of the momentum that the UAV already has gained. To make a controlled turn,

[5]The time duration of the active spoofing attack is the time for which GPS signals entirely control the UAV.

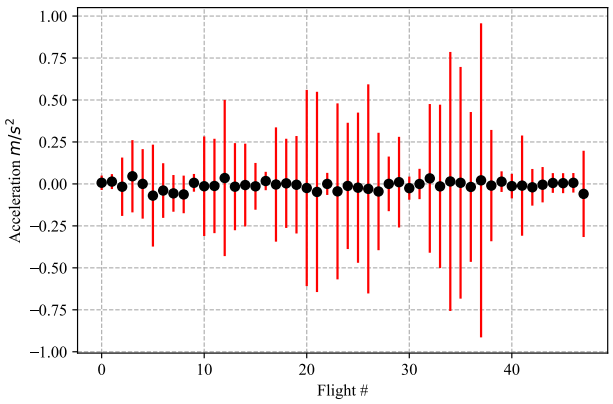

Figure 16: The average and standard deviation of instantaneous acceleration values of 48 flights with our velocity control algorithm. Flights with near-zero average acceleration achieve constant velocity.

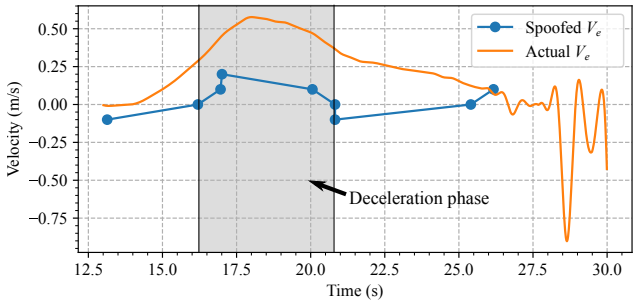

Figure 17: Effect of the proposed deceleration maneuver on the velocity of the target UAV. Notice the downward trend of the velocity post deceleration maneuver completion. An attacker can control the rate of deceleration by controlling the spoofed velocity.

it is vital to first null the UAV's velocity components v_e and v_n. In other words, the UAV must be stopped momentarily to enable sharp turns. The velocity control mechanism that we developed earlier was used to decelerate the UAV and reduce its speed to 0 m/s before changing the direction. To achieve this, we transmit a GPS signal that forces the UAV to decelerate for a longer duration, enough for the UAV to get its velocity to stay close to 0 m/s. Figure 17 shows the deceleration sequence that we developed. Notice how the spoofed V_e shifts between -0.1 m/s and 0.2 m/s. Figure 18 shows a representative flight where we employ the deceleration sequence to force the UAV to make a sharp 90° turn towards north.

4.4 Human-in-the-Loop (HITL) GPS Spoofing

Through all our experiments, we learned that precise real-time control requires us to control the acceleration of the

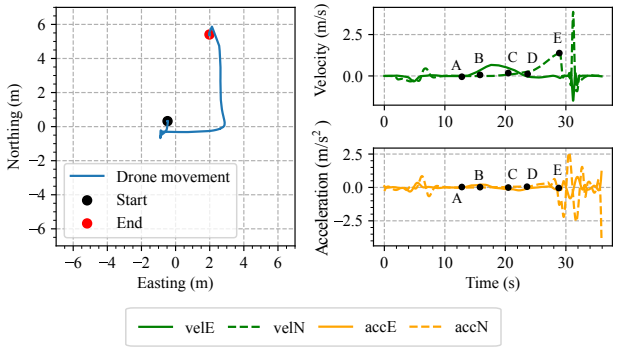

Figure 18: The target is executing a controlled sharp 90° turn. The attacker first forces the UAV to fly east. Then, at point B, the attacker starts the deceleration sequence, and at point D, the attacker forces the 90° turn.

target UAV. When the reaction time strategy fails, or the attacker cannot enumerate the reaction time, we need a feedback mechanism to observe the target's response to GPS spoofing and manipulate the spoofed trajectory accordingly. With the lessons learned through our spoofing experiments and the real-time capabilities of the GPS signal generator, we finally built and tested a Human-in-the-Loop (HITL) feedback system.

We have explored the possibility of using human intuition and knowledge of the target system to effectively control the target UAV, similar to a video game. To the best of our knowledge, such a system is a first-of-its-kind system designed to control a UAV via GPS spoofing. The HITL control system leverages RtGSG's ability to manipulate spoofed signals in real time through a human interaction device. In our experiments, it was a standard keyboard with arrow keys, but the device is interchangeable. This system relies on the attacker's observation of the UAV's movements and requires human intervention.

In our HITL system, the attacker uses the arrow keys to introduce a velocity vector in the form of $[v_e, v_n]$. This velocity vector governs the signal that the signal generator generates in real time. We modified the peripheral interface to reflect the operator's intentions and not directly apply the inputs to the spoofing signal. This interface shows the spoofed location and optionally the target's location if a system capable of tracking UAVs is available. The attacker then manipulates the velocity vectors based on their obtained understanding of the UAV's model, gaming skills, and intuition as to the UAV's possible reaction. Figure 19 shows one such flight where the operator takes manual control of the target UAVs and forces them to make controlled maneuvers, purely through GPS spoofing[6].

Limitations: The main objective of the HITL system is to provide control over the UAV just like a traditional remote-

[6]A video demonstration of this attack is available at https://www.youtube.com/watch?v=EtaQ_BQFn-M

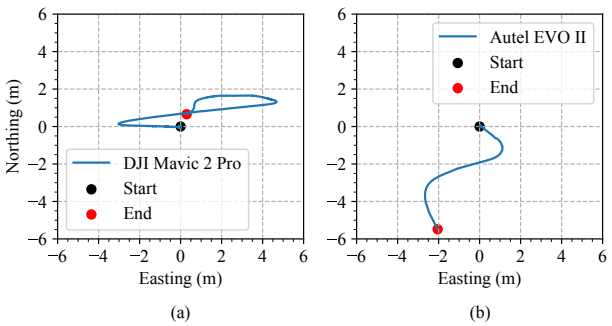

Figure 19: Plots (a - DJI Mavic 2 Pro and b - Autel EVO II) show the control using a Human-in-the-Loop control system.

controller, but through GPS spoofing. In the case of a regular controller, the control inputs directly actuate the motors. However, in the case of HITL GPS spoofing, the motors are actuated through the vehicle's attitude and position correction mechanism. This is the primary difference between controlling the UAV via GPS spoofing and controlling using a regular controller. Initially, one of the main challenges of operating the UAV indirectly through GPS spoofing is understanding that, e.g., spoofing signals that move right will result in the UAV drifting to the left. This requires the attacker to understand the motion dynamics of the UAV under GPS spoofing, and it requires training to maintain control of the UAV. Additionally, it also requires good hand-eye coordination

Furthermore, due to parallax misconceptions and response time delays, controlling the UAV via spoofed GPS signals is a much more challenging task than directly controlling the UAV through the original controller. Furthermore, the attacker needs to have a mechanism other than their own eyes to track and observe the target UAV. In other words, we believe that an automated closed-loop system would be an ideal way to execute a perfect takeover because such a system will overcome the discussed limitations.

5 Discussion

Forced Landing: The strategies we tested demonstrate an attacker's control over the horizontal position and velocity of the UAV. GPS spoofing can not be directly used to manipulate the height of the UAV and force it to land as majority of the UAVs rely on non-GPS sensors like rangefinders, downward-facing cameras, and barometers for measuring the altitude This poses a significant challenge because all these sensors are immune to GPS spoofing. For complete control of the UAV, the attacker should also land the UAV. As has been demonstrated earlier [34], an attacker can leverage terminal failsafes that UAVs implement to induce a forced landing. Some common events that trigger terminal failsafes are i) restricted zones and ii) EKF errors. Most consumer UAVs implement special geo-fencing around designated no-fly zones (NFZ) [21]. These restrictions prevent the UAV from taking off when the location is inside the no-fly zone. If the UAV accidentally enters a no-fly zone, the flight controller activates the terminal failsafe and lands after a warning. The operator can only control the UAV's horizontal position during this process.

The attacker's strategy thus is to spoof the GPS to a location inside the nearest no-fly zone in order to land the target. Suppose the no-fly location is more than 100 m away from the last spoofed location. In that case, the target temporarily loses the GPS lock but reacquires the attacker's signal within 20 s and initiates the landing sequence. In such a situation, time-to-first-fix (TTFF) for the onboard GPS receiver is typically between under 10 s as the receiver undergoes a warm start [73]. Even if the spoofed location is farther away and a warm start is not possible, the receiver performs a cold start, in which case the TTFF for a typical uBlox receiver 24-28 s [73]. Most of the UAVs that we tested were vulnerable to such an NFZ forced landing attack. Similarly, a UAV that implements an EKF failsafe can be forced to land by spoofing a motion that causes the position and velocity test ratios explained in Section 3.2, to cross the set threshold. It is important to note that not all manufacturers enforce such a failsafe, and hence such strategies to make the UAV land cannot be applied to every UAV. Table 1 summarizes the results of our experiments.

Limitations: Of the UAVs available to us, only the DJI Mavic 2 Pro, DJI Mavic Pro, and Autel EVO II allowed us to toggle downward vision sensors. This is an essential requirement for ensuring safety when testing inside the anechoic chamber, as often the UAV can behave erratically and crash. Other UAVs we tested were the DJI Mavic Mini and DJI Mavic Air 2, both of which primarily use vision sensors below a certain altitude. However, these models switch to GPS positioning above a certain altitude, which is greater than the height of the anechoic chamber available to us. Since it is illegal to transmit GPS signals over the air in the open, we could not test these UAVs. But based on our observations and the architecture of these UAVs (which are also from the market leader DJI), it is safe to say that the strategies that we proposed earlier apply to UAVs beyond the ones we tested.

In this experimental study, we target GPS. However, most modern UAVs are capable of multi-constellation localization. i.e., they use other satellite navigation systems like GLONASS, Galileo, and BeiDou. Such receivers can continue operations even when GPS is compromised. Multi-constellation receivers poses a challenge for successful spoofing and can be seen as a safeguard against GPS spoofing; however, these systems are known to be as vulnerable to signal spoofing attacks as GPS. Thus a multi-constellation GNSS signal generator can be used by an attacker to overcome this limitation as described for example in [57].

Table 1: A comparison of GPS takeover success and forced landing strategy success.

Model	GPS Takeover	Forced Landing	
		NFZ	EKF Failsafe
DJI Mavic Mini	Unable to test	✓	x
DJI Mavic Air 2	Unable to test	✓	x
DJI Mavic Pro	✓	✓	x
DJI Mavic 2 Pro	✓	✓	x
Autel Evo II	✓	x	x
3DR IRIS (sim)	✓	x	✓

Self-learning Feedback Mechanism: The self-learning feedback mechanism for controlling the spoofed trajectory requires access to precision UAV tracking equipment capable of UAV localization in real time, which serves as a source of the ground truth. In this system, Kalman-filter-based state estimation can be used to derive the target system's instantaneous acceleration and velocity through position information supplied by the tracking equipment. These values will be input to a predictive engine that can generate a trajectory capable of moving the UAV to the desired location. However, the attacker must consider the tolerable lag between the UAV's actual motion and the spoofed coordinates in this mechanism. This is especially applicable to UAVs that implement an EKF failsafe. Several works propose techniques to track UAVs, including acoustics sensors [12, 13], and works like [15, 28, 32] offer passive RADARs for tracking quadcopters. In [24], the authors suggest using a seeker UAV equipped with radios for target localization.

Impact of UAV Takeover: The UAV takeover strategies we propose in this work are capable of fine-grained direction and speed control of the target UAV. These capabilities allow an adversary to commandeer a UAV remotely with the intent of turning it rogue. Even a single rogue UAV, whether it is executing an autonomous mission by itself or as part of a swarm of UAVs, poses a significant security threat. We thus posit that GNSS as an attack vector needs to be considered systemically in conjunction with inertial, vision, and other sensors in future UAV security design.

On the other hand, precise control over a rogue UAV can help eliminate and investigate the threat, as the proposed techniques can be used defensively to get the rogue UAV to a safe location for further investigation.

Countermeasures: Current state-of-the-art countermeasures can be categorized as cryptographic solutions, physical and application layer solutions, and solutions that leverage IMUs for attack detection. UAVs typically have tight power and weight constraints, and hence the countermeasures should fit within these bounds. Thus, solutions that require minimal modifications to the existing infrastructure are essential.

Cryptographic solutions include techniques that introduce message encryption and authentication [14,29,43,45,76]. The underlying cryptographic primitives make it difficult for an attacker to synthesize GPS signals for arbitrary locations. However, these solutions are vulnerable to signal replay attacks. Furthermore, these countermeasures require a complete overhaul of the GPS ecosystem and are currently unavailable for evaluation. Additionally, they require high processing power, making it less practical to deploy on UAVs.

Physical and application layer countermeasures detect anomalies in RF properties like signal strength [3], auxiliary peaks [63], angle/direction of arrival [50,51], and validation of navigation messages like satellite ephemeris and timing information [19, 63]. Some countermeasures also leverage a multi-antenna and multi-receiver setup for attack detection [72]. An attacker can obfuscate physical properties through the careful generation of signals such that the signal's physical properties are within set thresholds. An attacker can also completely overshadow adversarial signals, thus burying the legitimate signal under the noise, thereby removing any auxiliary peaks. In addition, an adversary can use multiple antennas to mimic the angle and direction of legitimate signals.

There are proposals to use inertial sensor measurements to detect attacks on GPS through a comparison of independent inertial measurements and obtained GPS measurements [41, 44, 71]. Such solutions often use Kalman-filter-based sensor fusion algorithms already implemented in modern UAVs for state estimation. These techniques are effective against the proposed UAV takeover attack and can lead to attack detection. However, as shown in [53, 69, 79], an attacker can take over and defeat multi-sensor-fusion algorithms and inertial sensor solutions. Such a detection scheme can be used along with terminal failsafes [6], which force the UAV to abort any ongoing mission and either hold and hover at the current position or land. However, these techniques are limited to attack detection and do little in terms of attack mitigation and recovery. The receiver will also need to identify and attenuate adversarial signals in order to mitigate and recover from the attacks and ensure uninterrupted operations.

Several countermeasures use successive interference cancellation (SIC) technique [11, 26, 66] and antenna array processing techniques [17, 42, 48] to mitigate GPS spoofing attacks. Most of these mitigation techniques require high processing power or peripheral devices that make it impractical for implementation on UAVs. In [66], the authors have designed a solution specifically for UAVs. Such a countermeasure has the potential to recover from the GPS takeover attack proposed in this work. However, this solution is limited up to 15 dB of a spoofing signal's power advantage over legitimate signals beyond which the UAV cannot recover.

6 Related work

Several strategies have been proposed that demonstrate the ability of an attacker to assume control of a UAV via GPS spoofing. These works primarily focus on altering the motion of a UAV by transmitting fake GPS signals. In [57] the authors provide a taxonomy of hijacking consumer drones. However, the anti-drone hijacking strategies they propose are only capable of limited control over the target drone. Through a hard-spoofing attack, their strategies could divert a drone in a specific direction. Beyond that, the strategies proposed in this work do not facilitate the complete takeover of the target UAV as they lack post-takeover direction and velocity control ability. In [33, 40], through simulations, the authors demonstrated the effect of GPS spoofing on a cyber-physical system such as a UAV. Their approach forces the drone to accelerate in a particular direction by manipulating the GPS velocity in the opposite direction. In this approach, the direction of the motion and the UAV's acceleration was uncontrollable and unpredictable. As a result, such an approach cannot precisely control and maneuver the drone through GPS spoofing. In [47], the author demonstrated an attack that targets the follow-me feature specific to a DJI Phantom 3A. In this attack, the controller's mobile phone is targeted rather than the onboard GPS receiver. Since the attack only targets the "follow-me" flight mode, such an attack will not work against other completely autonomous flight modes. Similar works [7, 30, 67] demonstrate identical strategies that are restricted to a specific type of drone or a simulator environment, or that provide minimal true control over the target UAV.

There has been significant research on developing countermeasures to safeguard GPS receivers against spoofing attacks. Cryptographic solutions [14, 43, 45, 76] prevent attackers from generating counterfeit signals. However, they are still vulnerable to signal replay attacks and are not practical for deploying on a UAV because of additional processing power and key management requirements. The recently launched Galileo's Open Service Network Authentication [29] service that uses TESLA protocol for broadcast authentication is vulnerable to signal replay attacks. Other countermeasures like [10, 50–52, 72, 75] rely on peripheral hardware devices like multiple receivers and directional antennas, thus making them infeasible for integrating with UAVs. Another line of works [41, 71, 78] demonstrates the application of GPS/IMU sensor fusion to detect GPS spoofing attacks. However, as shown in [53, 79], it is possible to evade detection and defeat such multi-sensor fusion (MSF) algorithms. Shen et al. [69] analyzed the security guarantees of MSF algorithms implemented in terrestrial autonomous vehicles with the specific goal of forcing lane changes. Finally, [63] provides a technique that is based on signal processing and does not require additional hardware or cryptographic measures, but it does not protect against fine-grained spoofing of proximate locations as we have done in this work.

Several works have used non-GPS techniques for a hostile takeover of UAVs. These include the use of lasers to activate obstacle detection and avoidance systems [80], attacking the data-link between the radio-controller and the UAV [65] and [60] where the authors showcase the vulnerabilities present in a popular UAV platform from Parrot [1] as a result of a poorly configured wireless network.

7 Conclusion

In this work, we experimentally enumerated and validated various challenges in asserting complete control of a UAV through a GPS spoofing attack. We formulated requirements that constitute a complete takeover. To this extent, we designed, demonstrated, and evaluated strategies that enabled us to control the UAV's speed and direction in real time through well-timed GPS velocity manipulations that resulted in stable, predictable, and controlled flight. To facilitate real-time control, we designed and developed a real-time GPS signal generator capable of on-the-fly trajectory manipulation. We also designed a Human-in-the-Loop GPS spoofing system that can manually control a UAV's motion. Further, we discussed the possibility of incorporating an automatic self-learning feedback mechanism.

In conclusion, we show that even though the GPS receivers of COTS UAVs remain vulnerable to spoofing attacks, the combination of sensors incorporated in UAVs makes it extremely challenging for the attacker to translate a GPS spoofing attack into complete control over the UAV. We show in this work that — against conventional wisdom — only with a thorough study of a UAV's systemic behavior under GPS spoofing attacks and careful manipulation of the spoofing signals would it be possible to commandeer a UAV through GPS spoofing alone.

8 Acknowledgements

This work was partly supported by armasuisse Science and Technology (S+T), Thun, Switzerland. We also thank our shepherd Patrick Traynor and anonymous reviewers for their constructive comments and suggestions.

References

[1] Parrot. https://www.parrot.com/us.

[2] Space Segment. https://www.gps.gov/systems/gps/space/.

[3] Dennis M Akos. Who's afraid of the spoofer? gps/gnss spoofing detection via automatic gain control (agc). *NAVIGATION, Journal of the Institute of Navigation*, 2012.

[4] Alan (Big Al). Drone Payloads: Which Drone Can Carry The Most Weight? https://dronesvue.com/drone-payloads-which-drone-can-carry-the-most-weight/.

[5] ArduPilot. Arducopter. https://ardupilot.org/copter/.

[6] ArduPilot. EKF Failsafe. https://ardupilot.org/copter/docs/ekf-inav-failsafe.html.

[7] Sandra Pérez Arteaga, Luis Alberto Martínez Hernández, Gabriel Sánchez Pérez, Ana Lucila Sandoval Orozco, and Luis Javier García Villalba. Analysis of the GPS spoofing vulnerability in the drone 3DR solo. *IEEE Access*, 2019.

[8] Autel Robotics. Autel Robotics. https://www.autelrobotics.com/.

[9] Autel Robotics. User Manual - Autel EVO II Series. https://cdn.shopify.com/s/files/1/1538/0803/files/EVO_II_Series_User_Manual_En_Issue_Version.pdf.

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

[11] Ali Broumandan, Ali Jafarnia-Jahromi, and Gérard Lachapelle. Spoofing detection, classification and cancelation (SDCC) receiver architecture for a moving GNSS receiver. *GPS Solutions*, 2015.

[12] Joël Busset, Florian Perrodin, Peter Wellig, Beat Ott, Kurt Heutschi, Torben Rühl, and Thomas Nussbaumer. Detection and tracking of drones using advanced acoustic cameras. In *Unmanned/Unattended Sensors and Sensor Networks XI; and Advanced Free-Space Optical Communication Techniques and Applications*, 2015.

[13] Xianyu Chang, Chaoqun Yang, Junfeng Wu, Xiufang Shi, and Zhiguo Shi. A surveillance system for drone localization and tracking using acoustic arrays. In *Proceedings of the 2018 IEEE 10th Sensor Array and Multichannel Signal Processing Workshop (SAM)*, 2018.

[14] Xi-jun Cheng, Jiang-ning Xu, Ke-jin Cao, and Jie Wang. An authenticity verification scheme based on hidden messages for current civilian GPS signals. In *IEEE 2009 Fourth International Conference on Computer Sciences and Convergence Information Technology*, 2009.

[15] CRFS. RFeye DroneDefense. https://www.crfs.com/drone-detection/.

[16] Regulus Cyber. Ring - Revolutionary Next Generation Counter-UAS System. https://www.regulus.com.

[17] Saeed Daneshmand, Ali Jafarnia-Jahromi, Ali Broumandan, and Gérard Lachapelle. A GNSS structural interference mitigation technique using antenna array processing. In *2014 IEEE 8th Sensor Array and Multichannel Signal Processing Workshop (SAM)*, 2014.

[18] Isabel Debre. Saudi Arabian oil facility struck in drone attack, 2021. https://www.bbc.com/news/world-middle-east-60082786.

[19] DHS. Positioning, Navigation, and Timing (PNT) Program. https://www.dhs.gov/science-and-technology/pnt-program.

[20] DJI. DJI. https://www.dji.com.

[21] DJI. Geo Zone Map - Fly Safe - DJI. https://www.dji.com/flysafe/geo-map.

[22] DJI. Mavic 2 - Product Information - DJI. https://www.dji.com/mavic-2/info.

[23] DJI. Mavic 3 - Specs - DJI. https://www.dji.com/mavic-3/specs.

[24] Louis Dressel and Mykel J Kochenderfer. Hunting drones with other drones: Tracking a moving radio target. In *International Conference on Robotics and Automation (ICRA)*, 2019.

[25] Takuji Ebinuma. Software-Defined GPS Signal Simulator, 2015. https://github.com/osqzss/gps-sdr-sim.

[26] Manuel Eichelberger, Ferdinand Von Hagen, and Roger Wattenhofer. A Spoof-Proof GPS Receiver. In *19th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2020.

[27] ETS-Lindgren. Broadband Mini-Bicon Antenna - ETS-Lindgren. https://www.ets-lindgren.com/products/antennas/broadband-min-bicon-antennas/4005/400502?page=Products-Item-Page.

[28] Gao Fang, Jianxin Yi, Xianrong Wan, Yuqi Liu, and Hengyu Ke. Experimental research of multistatic passive radar with a single antenna for drone detection. *IEEE Access*, 2018.

[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*, 2016.

[30] João Gaspar, Renato Ferreira, Pedro Sebastião, and Nuno Souto. Capture of UAVs through gps spoofing using low-cost SDR platforms. *Capture of UAVs through GPS spoofing using low-cost SDR platforms*, 2020.

[31] Werner Gurtner and Lou Estey. Rinex-the receiver independent exchange format-version 3.00. *Astronomical Institute, University of Bern and UNAVCO, Bolulder, Colorado.*, 2007.

[32] İsmail Güvenç, Ozgur Ozdemir, Yavuz Yapici, Hani Mehrpouyan, and David Matolak. Detection, localization, and tracking of unauthorized uas and jammers. In *IEEE/AIAA 36th Digital Avionics Systems Conference (DASC)*, 2017.

[33] Daojing He, Yinrong Qiao, Shiqing Chen, Xiao Du, Wenjie Chen, Sencun Zhu, and Mohsen Guizani. A friendly and low-cost technique for capturing non-cooperative civilian unmanned aerial vehicles. *IEEE Network*, 2018.

[34] Ling Huang and Qing Yang. Low-cost GPS simulator GPS spoofing by SDR. In *Proc. DEFCON*, 2015.

[35] NaturalPoint Inc. Motion Capture Systems. https://www.optitrack.com/.

[36] Drone Industry Insights. Top 10 Drone Manufacturer's Market Shares in the US. https://droneii.com/product/drone-manufacturers-ranking.

[37] National Instruments. Ettus Research. https://www.ettus.com/products/.

[38] Mordor Intelligence. Drones Market - Growth, Trends, COVID-19 Impact, And Forecasts (2022 - 2027), 2022. https://www.mordorintelligence.com/industry-reports/drones-market.

[39] Xichen Jiang, Jiangmeng Zhang, Brian J Harding, Jonathan J Makela, Alejandro D Domı, et al. Spoofing GPS receiver clock offset of phasor measurement units. *IEEE Transactions on Power Systems*, 2013.

[40] 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*, 2014.

[41] Samer Khanafseh, Naeem Roshan, Steven Langel, Fang-Cheng Chan, Mathieu Joerger, and Boris Pervan. Gps spoofing detection using raim with ins coupling. In *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 2014.

[42] Andriy Konovaltsev, Stefano Caizzone, Manuel Cuntz, and Michael Meurer. Autonomous spoofing detection and mitigation with a miniaturized adaptive antenna array. In *Proceedings of the 27th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2014)*, 2014.

[43] Markus G Kuhn. An asymmetric security mechanism for navigation signals. In *International Workshop on Information Hiding*, 2004.

[44] 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*, 2019.

[45] Sherman C Lo and Per K Enge. Authenticating aviation augmentation system broadcasts. In *Proceedings of the IEEE/ION Position, Location and Navigation Symposium*, 2010.

[46] Lime Microsystems Ltd. LimeSDR. https://limemicro.com/products/boards/limesdr/.

[47] Aaron Luo. Drones Hijacking - multi-dimensional attack vectors and countermeasures, 2016. https://www.youtube.com/watch?v=u9nFdOvA8eI&ab_channel=SecurityHub.

[48] Jaroslaw Magiera and Ryszard Katulski. Detection and mitigation of GPS spoofing based on antenna array processing. *Journal of applied research and technology*, 2015.

[49] Jonathan Marcus. Yemen rebel attack on UAE throws challenge to the region, 2022. https://www.bbc.com/news/world-middle-east-60082786.

[50] Charles E McDowell. GPS spoofer and repeater mitigation system using digital spatial nulling, 2007. US Patent 7,250,903.

[51] Michael Meurer, Andriy Konovaltsev, Manuel Appel, and Manuel Cuntz. Direction-of-Arrival Assisted Sequential Spoofing Detection and Mitigation. In *Proceedings of the 2016 International Technical Meeting of The Institute of Navigation, Monterey, California*, 2016.

[52] Paul Y Montgomery. Receiver-autonomous spoofing detection: Experimental results of a multi-antenna receiver defense against a portable civil GPS spoofer. In *Proceedings of the Radionavigation Laboratory Conference*, 2011.

[53] Sashank Narain, Aanjhan Ranganathan, and Guevara Noubir. Security of GPS/INS based on-road location tracking systems. In *IEEE Symposium on Security and Privacy (S&P)*, 2019.

[54] BBC News. Heathrow airport: Drone sighting halts departures. 2019. https://www.bbc.com/news/uk-46803713.

[55] 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/.

[56] 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.

[57] 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)*, 2019.

[58] Ophir RF Inc. Ophir RF Model 5293. https://ophirrf.com/wp-content/uploads/2015/09/5293-2.pdf.

[59] Josh Petri. How Hackers Can Take Over Your Car's GPS, 2019. https://www.bloomberg.com/news/articles/2019-06-19/threat-of-gps-spoofing-for-autonomous-cars-seen-as-overblown.

[60] Johann-Sebastian Pleban, Ricardo Band, and Reiner Creutzburg. Hacking and securing the AR. Drone 2.0 quadcopter: investigations for improving the security of a toy. In *Mobile Devices and Multimedia: Enabling Technologies, Algorithms, and Applications*, 2014.

[61] Associated Press. UAE Bans Flying of Recreational Drones After Fatal Attack. https://www.voanews.com/a/uae-bans-flying-of-recreational-drones-after-fatal-attack/6408720.html.

[62] Mark L Psiaki and Todd E Humphreys. GNSS spoofing and detection. *Proceedings of the IEEE*, 2016.

[63] Aanjhan Ranganathan, Hildur Ólafsdóttir, and Srdjan Capkun. SPREE: A spoofing resistant GPS receiver. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*, 2016.

[64] Open Robotics. Gazebo. http://gazebosim.org/.

[65] Nils Miro Rodday, Ricardo de O Schmidt, and Aiko Pras. Exploring security vulnerabilities of unmanned aerial vehicles. In *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2016.

[66] 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.

[67] Seong-Hun Seo, Byung-Hyun Lee, Sung-Hyuck Im, and Gyu-In Jee. Effect of spoofing on unmanned aerial vehicle using counterfeited GPS signal. *Journal of Positioning, Navigation, and Timing*, 2015.

[68] Samira Shackle. The mystery of the Gatwick drone. 2020. https://www.theguardian.com/uk-news/2020/dec/01/the-mystery-of-the-gatwick-drone.

[69] Junjie Shen, Jun Yeon Won, Zeyuan Chen, and Qi Alfred Chen. Drift with devil: Security of multi-sensor fusion based localization in high-level autonomous driving under {GPS} spoofing. In *29th USENIX Security Symposium (USENIX Security 20)*, 2020.

[70] Daniel P Shepard, Jahshan A Bhatti, and Todd E Humphreys. Drone hack: Spoofing attack demonstration on a civilian unmanned aerial vehicle. *GPS World*, 2012.

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

[72] 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*, 2011.

[73] u-Blox. NEO-M8 DataSheet. https://www.u-blox.com/en/ubx-viewer/view/NEO-M8-FW3_DataSheet_UBX-15031086.pdf.

[74] James Vincent. Recreational drones banned in United Arab Emirates after oil facility attack, 2022. https://www.theverge.com/2022/1/24/22898614/united-arab-emirates-uae-ban-recreational-drone-attack.

[75] Jon S Warner and Roger G Johnston. GPS spoofing countermeasures. *Homeland Security Journal*, 2003.

[76] Kyle Wesson, Mark Rothlisberger, and Todd Humphreys. Practical cryptographic civil GPS signal authentication. *NAVIGATION: Journal of the Institute of Navigation*, 2012.

[77] Michael G Wing, Aaron Eklund, and Loren D Kellogg. Consumer-grade global positioning system (GPS) accuracy and reliability. *Journal of forestry*, 2005.

[78] Tao Zhang and Quanyan Zhu. Strategic defense against deceptive civilian GPS spoofing of unmanned aerial vehicles. In *Proceedings of the International Conference on Decision and Game Theory for Security*, 2017.

[79] Zhongshun Zhang, Lifeng Zhou, and Pratap Tokekar. Strategies to design signals to spoof Kalman filter. In *Annual American Control Conference (ACC)*, 2018.

[80] Ce Zhou, Qiben Yan, Yan Shi, and Lichao Sun. DoubleStar: Long-range attack towards depth estimation based obstacle avoidance in autonomous systems. *arXiv preprint arXiv:2110.03154*, 2021.