

Drones Security and Privacy: Attacks to Sensors

CPS and IoT Security

Alessandro Brighente

*Master Degree in
Cybersecurity*



UNIVERSITÀ
DEGLI STUDI
DI PADOVA



SPRITZ
SECURITY & PRIVACY
RESEARCH GROUP

Unmanned Aerial Vehicles



- Unmanned Aerial Vehicles (UAVs) are flying devices without any crew or pilot on board
- They come in different shapes and size, and can reach different altitudes



Fixed-wing rotor




Quadcopter



Tricopter

- UAVs may be operated by pilots located on the ground
- Remotely Piloted Aircraft Systems (RPAS)
- Formally, “a set of configurable elements consisting of a remotely piloted aircraft, its control station, the command and control links and any other system elements required during flight operation”



Operating Mode (2)



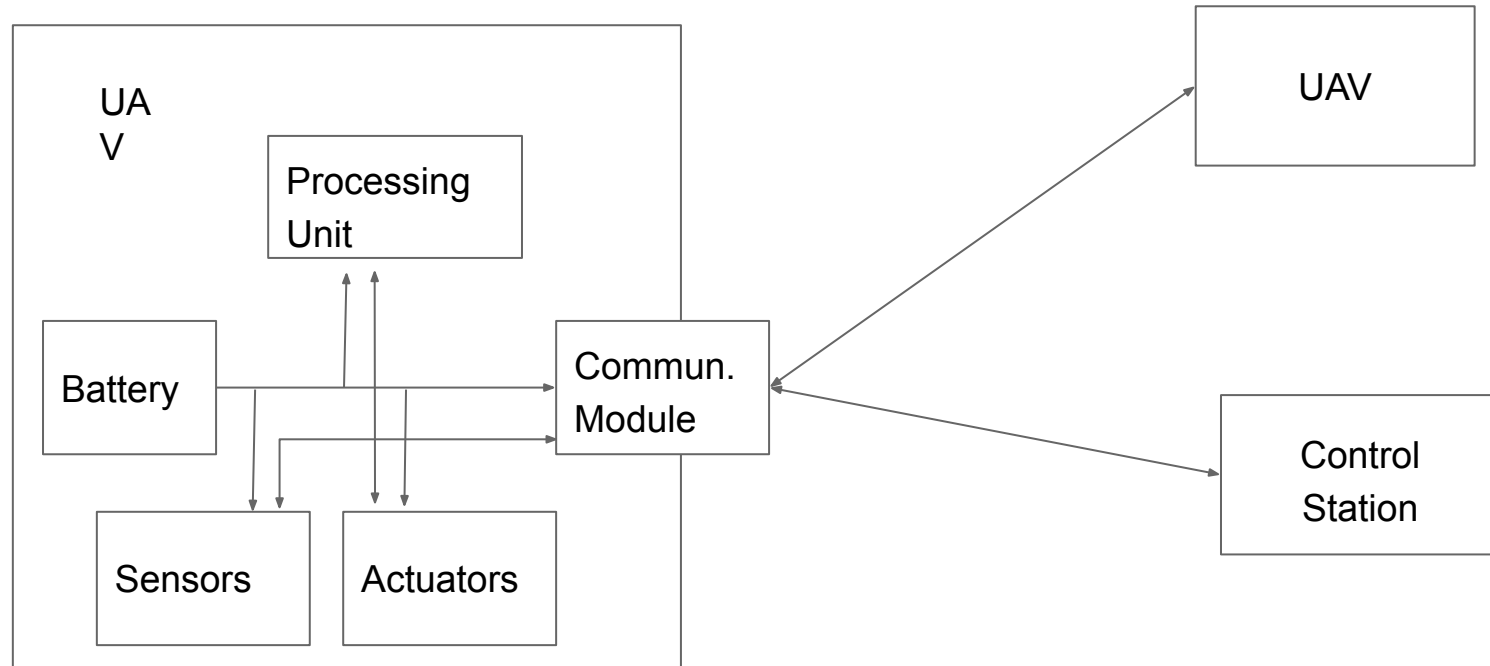
- UAVs may be autonomous and coordinate with other drones to deliver a common objective
- Fleet of UAVs
- Need to communicate with other fleets to avoid crashes



Structure of a Drone



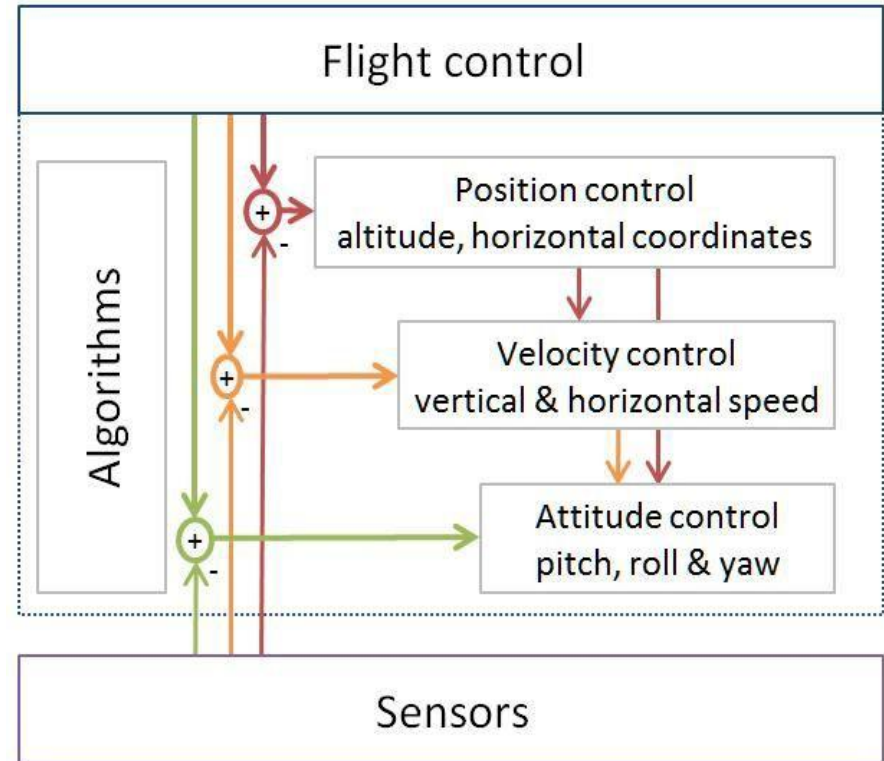
- Multiple modules to acquire and process data
- Communication module as enabler



Operating Principle



- UAVs employ control loops possibly with sensor feedback



By Maxorazon - Own work, CC BY-SA 4.0,
<https://commons.wikimedia.org/w/index.php?curid=46770558>



- UAVs have been firstly proposed for military applications
- Seventeen countries armed UAVs, and more than 100 countries use UAVs in military capacity
- Civil applications are instead more recent
- Thanks to their high adaptability, UAVs find applications in many different fields
 - Infrastructure
 - Transport
 - Media and Entertainment
 - Telecommunication
 - Agriculture
 - Search and Rescue



- A drone ecosystem is composed by six unique targets:
 1. Drone hardware: CPU, sensors, firmware
 2. Drone chassis and package: all non-electronic devices
 3. Ground control station: may be fixed or mobile
 4. First-Person View (FPV) channel: control channel via common communication protocols
 5. Pilot: person remotely controlling the drone
 6. Cloud services: some drone send telemetry of flight information to a cloud server if needed



- Attacker aim: disrupt the legitimate task of flying a drone
- Attackers may be both civilian or military
- Although military may have sophisticated equipments (e.g., cannons or predator birds), a civilian may still be able to impact on the drone
- We consider a civilian that has access to equipments that can be purchased at a moderate cost
 - Software-Defined Radio
 - Computer
 - Commercial lasers
 - Butterfly net
 - Magnets



- An attacker may be of three types:
 - With direct physical access → access the drone or the GCS to modify the firmware or replace hardware parts
 - Physically proximate → send, modify, and replay radio transmissions in order to hijack a drone
 - Distant adversary → resides on the Internet and applies attacks against servers, drones, the GCS
- The impact of the attack can be measured according to the CIA triad
 - Confidentiality: attack that reveals information about drone, pilot, telemetry or collected data
 - Integrity: attack that modifies the information delivered to or collected by the drone, GCS, cloud server or pilot
 - Availability: attack that causes the pilot to lose control of the drone due to forced landing, crash or hijacking



- Find vulnerabilities in the devices or software used by the drone
- Attack to the compass
 - Use a magnetic field to fool the compass and either hijack or prevent take off
- Attack on the stabilizing algorithm via camera sensor
 - Drones use a camera to collect images and stabilize the flight
 - To detect movements, UAVs use optical flow
 - Features can be modified by projecting laser beams on the ground, or projecting images on the ground
 - The drone will follow the projected features



- Optical flow is the pattern of apparent motion of objects, surfaces, and edges (basically, everything from which we can extract simple features) in a visual scene caused by the relative motion between an observer and the scene
- Optical flow uses a sequence of successive images to allow for the estimation of motion
- It tries to calculate the motion between two image frames which are taken at a predefined rate at every voxel position

- UAVs use a downward facing optical flow camera
- They collect images and use the estimation to stabilize the flight
- It can detect whether the drone is drifting by comparing successive frames of the ground below
- The sensor system will attempt to infer if the ground plane image has moved by a relative offset $(\Delta x, \Delta y)$
- If so, the system assume that the ground is stationary and the drone has drifted





- Optical flow first requires a feature detection algorithm to identify regions of the ground plane easy to track
- These features are then provided as input to the optical flow algorithm
- It then identifies the location of these features in successive images and uses the difference to compute the displacement
- Classic configuration: use the Shi-Tomasi corner detection algorithm for feature detection and the Lucas-Kanade method to compute the optical flow



- Basic idea: use the derivative of the image to detect whether there is a sudden change of color in one direction
- If only in direction, then an edge is detected
- If in two direction, then a corner is detected
- It is generally used because efficient, compared to more sophisticated but slower feature detectors (SURF, SIFT)



- Assumes that the difference between two consecutive frames is small and approximately constant within some neighbourhood
- This assumption is arranged for UAVs considering a combination of sufficiently high framerate and sufficiently low resolution
- For each corner pixel p returned by the feature detector, let us define a local window of n neighbourhood pixels for $i \in \{1, \dots, n\}$
- For each i let $I_x(q_i)V_x + I_y(q_i)V_y = -I_t(q_i)$

where I denote the image matrix, I_x denotes partial derivative, and V the velocity or motion along a coordinate \rightarrow estimate V via Least Squares



- Build system

$$I_x(q_1)V_x + I_y(q_1)V_y = -I_t(q_1)$$

$$I_x(q_2)V_x + I_y(q_2)V_y = -I_t(q_2)$$

⋮

$$I_x(q_n)V_x + I_y(q_n)V_y = -I_t(q_n)$$

- Matrix form $Av = b$



- The system has more equations than unknowns → overdetermined
- Compromise solution via the least squares principle

$$A^T A v = A^T b \longrightarrow v = (A^T A)^{-1} A^T b$$

- Computes

$$\begin{bmatrix} V_x \\ V_y \end{bmatrix} = \begin{bmatrix} \sum_i I_x(q_i)^2 & \sum_i I_x(q_i)I_y(q_i) \\ \sum_i I_y(q_i)I_x(q_i) & \sum_i I_y(q_i)^2 \end{bmatrix}^{-1} \begin{bmatrix} -\sum_i I_x(q_i)I_t(q_i) \\ -\sum_i I_y(q_i)I_t(q_i) \end{bmatrix}$$



- A successful attack to sensors needs three requirements
- Environment influence: the adversary can alter the physical phenomenon that the system measures (alter pixel values)
- Plausible input: create an input to the sensor that will actually be used by the system as valid input
- Meaningful response: the attacker can induce a behavior on the UAV representing the control the attacker has over it



- The adversary must be able to alter or obscure the ground plane to alter the pixel values reported by the optical flow camera
- Onerous example: cover part of the area with feature-rich pattern
- More concrete example: project features on the ground via laser/light beams
- In the first case, the attacker can use a battery operated projector loading images e.g., from a USB stick
- In the second case, use an array of laser elements to create features



- The basic idea of optical flow is to instruct the UAV to compensate a drift by moving of the same amount in the opposite direction
- The system assumes the image on the ground to be stationary, so it interprets feature motion along a vector as a movement in the opposite direction
- By moving the ground feature, the attacker can control the movement of the drone



- The attacker's goal is to project a sharp gradient onto the ground so that the feature detection algorithm will pick up the light as a corner
- Since the Lucas-Kanade-based optical flow computes a final displacement based on the average displacement of each feature, the attacker needs to generate a large number of corners
- In practice, the attacker can simply sweep their projected light across the ground plane

- Most UAVs are enabled with the Return to Home (RTH) functionality
- The UAV records its location of departure, and in case of emergency it will automatically fly back to home
- Emergency includes low battery or loss of the control signal for more than 3 seconds

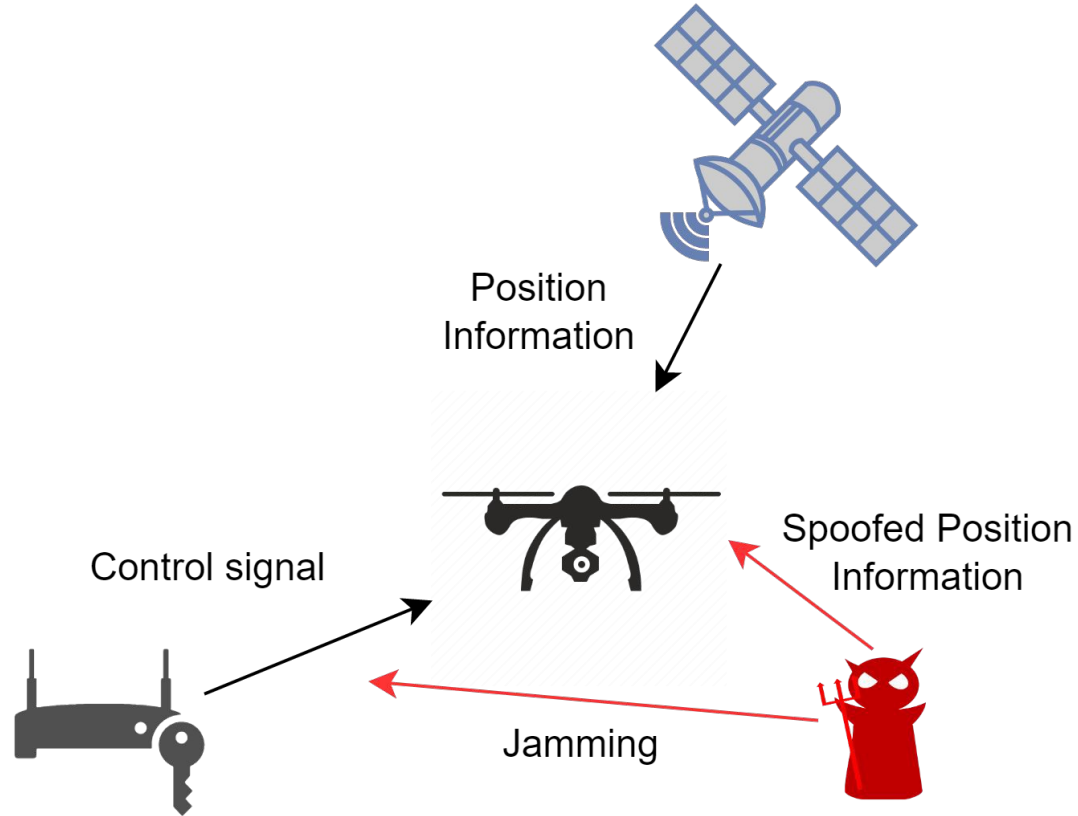


From DJI's RTH guide



- Most UAVs rely on the Global Navigation Satellite System (GNSS) for positioning (e.g., Europe's Galileo)
- However, this system has the main issue of not being authenticated
- We can hence spoof the GNSS signal and control the location of the drone
- However, the drone is controlled by its pilot

Capturing Drones





- When jamming the signal for enough time, the drone enters the RTH mode and starts to fly back to the home location
- With the help of the GPS module and the on-board compass, it can derive a trajectory based on its current location S and the home location H
- The attacker aims at sending spoofed location information to capture the drone in the minimum time by bringing it to a controlled location D




- To achieve the shortest capturing time, the attacker may want the UAV to fly along the straight line current-desired location (SD)
- However, the drone has a pre-recorded H and will try to return there
- If only one could spoof an arbitrary location to the UAV, then the UAV will change its direction immediately to any desired angle
- For instance, given a location S' such that $S'H \parallel SD$, then the UAV will fly to D along SD
- However, GPS accepts only physics-reasonable coordinates

Greedy Algorithm for UAV Capture



- Constraint: the maximum location change speed of GPS spoofing is V_s
- Idea: to minimize the capture time, the UAV trajectory shall be spoofed such that it points towards D as soon as possible





- Loiter mode: the UAV tries to hold its position S by correcting any drift. We exploit this feature to have the drone gradually moving to D
- An ideal spoofing trajectory is the one starting from the horizontal line a (previous figure), where the arriving point is exactly above H
- Besides spoofing, the UAV will always point towards the direction of H tanks to its compass
- Whenever it believes it is on the left side of $S1'$, it moves closer to H horizontally



- After passing $S1$, the UAV direction will turn left until it points precisely at D
- Denoting $S2'$ as the turning point in the spoofed trajectory and its corresponding point in the actual trajectory as $S2$, $S2'H \parallel S2D$
- The UAV, convinced to be taking the spoofed trajectory $S2'H$, will fly along the straight line $S2D$ until it reaches D



- How to determine the spoofed trajectory $S1'S2'$ such that its actual trajectory becomes $S1S2$ and $S2'H$ becomes parallel with $S2D$ at the earliest time?
- Greedy algorithm to compute $S2'$
- At the spoofed location $S1'$, the next spoofed point Sx' is such that $S1Sx'$ is perpendicular to $S1D$
- If $SxD \parallel Sx'H$, done
- Otherwise, repeat
- See python code example



- GPS systems may occasionally drop the signal or suffer from glitches, i.e., provide significantly inaccurate position information
- The **GPS failsafe** is a mechanism to provide safety to the drone in case of GPS signal loss or glitches
- In this case, the drone may either land or switch to a manual control
- To protect against glitches, the drone has a short memory of GPS position and compares the new measurement with the old one



- Extended Kalman Filter implementation
 - When a new GPS measurement is received, compare it to a position predicted based on IMU measurements
 - If the difference is above a threshold, discard it
 - While GPS measurements are not used, we define an uncertainty radius around the predicted point
 - If subsequent measurements fall in the circle, they will be used and the circle is reset to the minimum radius



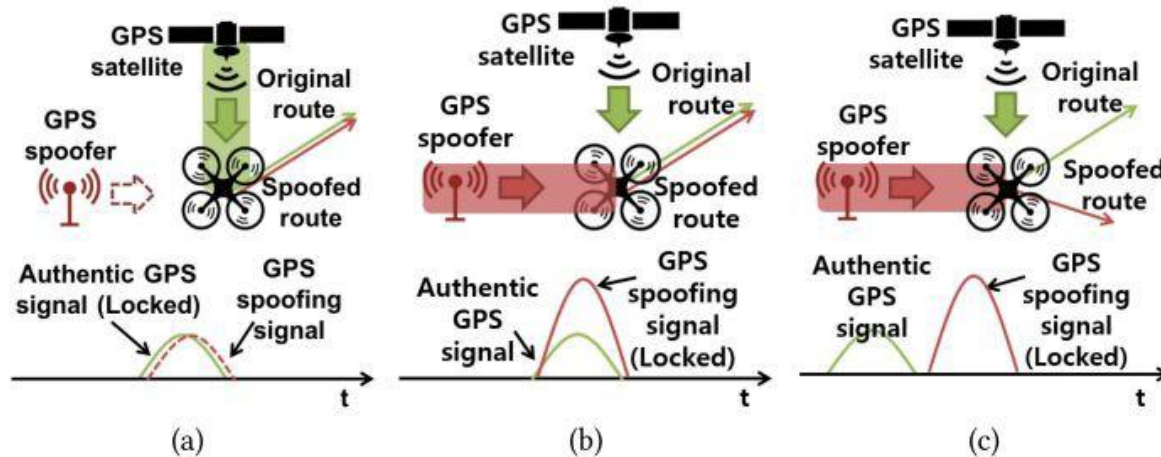
- We can classify consumer grade drones based on their behavior after failsafe mode, i.e., locking GPS again
- First type: switch from the failsafe mode to positioning mode which utilizes GPS
- Second type: resume autopilot just before loss-of-lock
- Third type: maintain GPS failsafe even though GPS is available and wait until a pilot gives a new command



- We want to develop an anti-drone system such that, as soon as we detect the drone, we do not simply interrupt the control channel and have the drone to hover over a sensitive area
- To carry out a mission that evades RC jamming by anti-drone solutions, attackers will operate drones in autopilot mode based on GPS and will not rely on a remote controller
- Therefore, spoofing GPS might be a solution to safely remove drones from an area they should not be into

Soft GPS Spoofing

- The spoofing signal is aligned with the authentic GPS signal
- The operation is not interrupted and the victim gradually locks to the spoofed GPS signals in three steps





Hard GPS Spoofing

- Soft spoofing has some requirements to satisfy to avoid losing the lock
- When one of the requirements is not satisfied, we call this signal as hard GPS spoofing
- It initially acts like a jamming signal and the victim may lose its lock on the authentic signal
- The spoofing signal is stronger, thus the victim will reconnect to the spoofed signal



- It deals with type 1 drones
- The drone is trying to maintain its original position
- The attacker spoofs the target drone GPS position as if the drone is moving in a certain direction
- The target drone generates speed in the opposite direction, so the drone moves in that direction in the real world



- It deals with type 2 drones
- Based on the drone's characteristic that control their body according to their path-following algorithm (e.g., in RTH)
- If the GPS position is manipulated as the drone deviates from the path, then it will move in a different direction from the original direction to return to the track
- The hijacker can determine the hijacking direction and calculate the corresponding fake location



- It deals with type 3 drones
- Since they need to wait for a pilot command, manipulation of GPS signal is not effective
- Therefore, we want to hijack them without losing the connection with the GPS (soft GPS spoofing)
- However, it might be complicated to move the drone if it uses a combination of GPS and IMU sensors to determine its location




- Case study against DJI Phantom 3 and 4
- GPS position and velocity of drones can be manipulated as intended by a hard GPS spoofed signal
- By continuously transmitting GPS signals stronger than the legitimate ones, the drone locks to the attacker
- However, we notice that there exists an angular error between the expected direction and the measured path
- [Link](#) to demonstrations

Case Study B: Parrot Bebop 2



- Parrot Bebop 2 provides flight plan mode
- To control the path-following algorithm, we manipulate the GPS location of the drone out of its track via hard GPS spoofing while operating in flight plan mode
- Result: its moving direction is always the direction from its current location to the next waypoint





Safe hijack procedure:

- The attacker finds the location of the waypoint p_{waypoint} by manipulating GPS and bring it out of track (see previous figure)
- As the new direction will point to the waypoint, we can draw a straight line connection the drone movement and intersect it with the original one to find the waypoint
- The attacker then determines the fake location and generate a fake GPS signal accordingly



- Assume the attacker has a target position p_{target}
- The fake location should be a point on the line passing through the waypoint in the intended hijacking direction

$a = p_{\text{waypoint}} + k (p_{\text{target}} - p_{\text{init}})$, where $k < 0$ is an adjustable parameter



- 3DR Solo does not have a fallback mechanism when losing connection with the GPS
- It is complicated to hijack, because it uses both GPS and IMU for localization
- It has a complicated path-following algorithm that uses Intermediate Target Position (ITP), necessitating moving the spoofed GPS location accordingly to safely control the behavior of the target drone
- Good case study for strategy C!



- The drone uses an Extended Kalman Filter algorithm to estimate various parameters, such as velocity, position, and magnetic field by fusing the IMU's output with GPS measurements
- The EKF algorithm starts by predicting its position and velocity by using the IMU's output only
- IMUs have inherent errors → periodically compare its predictions with the GPS values and correct based on errors

ALGORITHM 1: EKF failure detection algorithm

```
// ekf_check() is called at 10 Hz by ArduCopter scheduler
```

```
1 Function ekf_check()  
2   if motors are stopped then  
3     bad_variance ← False  
4     fail_count ← 0  
5     return  
6   get innovVelSumSq from the EKF  
7   get varVelSum from the EKF  
8   velVar ← sqrt(innovVelSumSq / varVelSum)  
9   if velVar ≥ 0.8 then  
10    if bad_variance is not True then  
11      fail_count ← fail_count + 1  
12      if fail_count ≥ 10 then  
13        fail_count ← 10  
14        bad_variance ← True  
15        change its mode to the EKF fail-safe mode  
16  else  
17    if fail_count > 0 then  
18      fail_count ← fail_count - 1  
19      if bad_variance is True and fail_count == 0 then  
20        bad_variance ← False
```

Called every 100 ms

Sum of innovations along axis




- During flight the drone might deviate from the track because of external influences such as wind
- The path following algorithm is what prevent the drone from completely deviating from its mission
- The ArduCopter path following algorithm is based on the ITP position
- ArduCopter periodically advances the ITP along the track in small increments and causes the drone body to move to the ITP

Path Following Algorithm



ALGORITHM 2: ArduCopter path-following algorithm

```
// The original name of the function is advance_wp_target_along_track()
// advance_target() is called at 400 Hz by the ArduCopter scheduler
1 Function advance_target()
  // ITP is initialized with the starting point of the track
2  get track from the mission uploaded by a user
3  location ← the coordinates of the current location from the EKF
4  distance ← the distance from the location to the track
5  if distance is less than 13 m then /* Zone 0 */
6    | ITP advances slightly forward along the track;
7  else
8    perpendicular_foot ← the coordinates of perpendicular foot from the location to the track
9    if ITP is closer to the end of the track than perpendicular_foot then /* Zone 1 */
10   | ITP stays
11   else
12     if distance / (the distance between location and ITP) > 0.98 then /* Zone 2 */
13     | ITP advances slightly forward along the track
14     else /* Zone 3 */
15     | ITP advances slightly forward along the track until the drone's speed along the track
16     | exceeds 1 m/s
17     | ITP halts when the drone's speed along the track exceeds 1 m/s
17 | controls motors to move the drone body to ITP
```



- ITP remains unchanged if the drone is in Zone 1
- Advances slightly forward along the track if in Zone 2 or 3
- From the attacker's point of view, the ITP can be estimated to be the physical location of the drone immediately prior to GPS spoofing



- If the drone mistakenly determines that it has deviated from the track owing to GPS spoofing, then it will move in the direction away from the manipulated location to the ITP
- The path following algorithm uses a *leash length*, i.e., a minimum distance from the current position to the main path
- The attacker can exploit this distance to hijack the drone's location

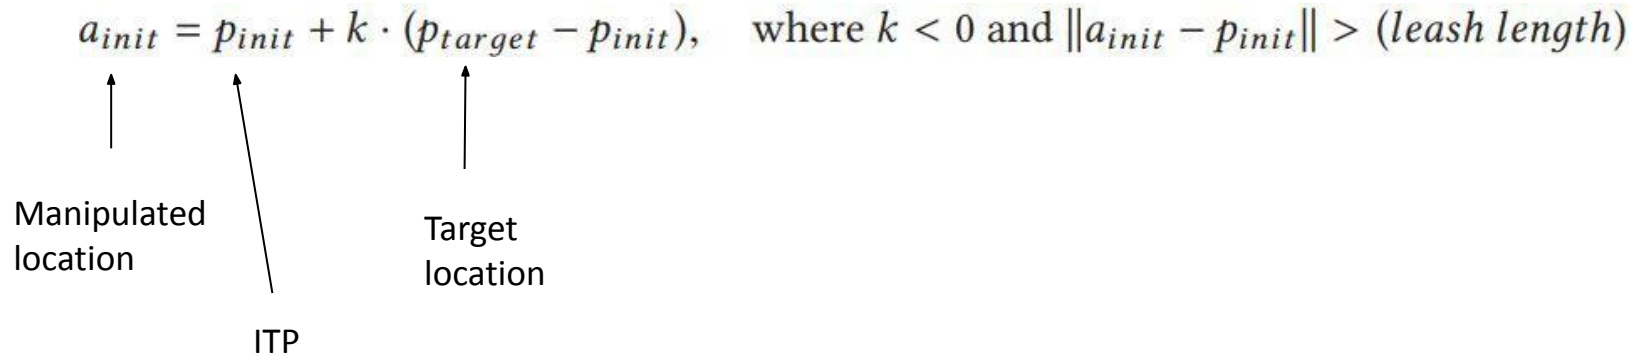


- The ITP is fixed or only slightly changed when the GPS location deviates from the main track by more than the leash length
- We can hence say that the physical location of the drone immediately before hijacking is an approximation of the updated ITP
- The attacker hence derive the possible coordinates of the initial fake location from the approximated ITP and the intended hijack direction



- We assume the attacker begins hijacking at the physical location $p_{init} = (p_1, p_2, p_3)$
- Objective: move the drone to p_{target}
- Hijacking direction = line between p_{init} and p_{target}

- The drone will try to move to the ITP (approximately p_{init}) if its GPS location deviates from the track more than the leash length
- The direction from the initial fake location $a_{init} = (a_1, a_2, a_3)$ to the ITP should be the same as the hijacking direction

$$a_{init} = p_{init} + k \cdot (p_{target} - p_{init}), \quad \text{where } k < 0 \text{ and } \|a_{init} - p_{init}\| > (\text{leash length})$$


Manipulated location

ITP

Target location



- The equation is the vector form of the parametric equation of a line
- Thus, a_{init} lies on the line that passes from pinot in the direction from target to pinot
- In addition, the distance between pinot and a_{init} should be higher than the leash length



- The `fail_count` value will increase by one if the GPS location of the drone jumps to a_{init} because the GPS velocity and that measured by IMUs is not consistent
- Therefore, the GPS velocity must be changed adaptively to be similar to the motion of the drone body to avoid triggering the failsafe mode
- It makes the GPS location approach the drone's original track to within the leash length, but we can prevent a change in the ITP by manipulating the GPS location to a_{init} again



- The hijacker repeats the process until the drone reaches the safe target location
- The EKF failure count increases by one for each jump, but it will decrease to zero if the GPS velocity is consistent of that of the IMUs
- We define as t' the current time, and as Δ the GPS update period
- We update the current fake location as

$$a_{t=t'} = a_{t=(t'-\Delta)} + (p_{t=t'} - p_{t=(t'-\Delta)})$$




- A drone needs a controller able to adjust the horizontal speed at a rotor level
- Indeed multiple rotors are not always exactly the same and the center of mass cannot always be ensured
- A flight attitude controller is hence implemented at a software level
- Compute the proper control signal for multiple rotors based on the data from Inertial Measurement Units (IMU), including gyroscopes



- An IMU is an electronic device using a combination of accelerometers, gyroscopes, and sometimes magnetometers to measure a body's specific force
- In a UAV, an IMU measures the orientation, rotation, and acceleration
- Micro-Electro-Mechanical Systems (MEMS) gyroscopes are used to make flight control modules small

- The principle of the MEMS gyroscope is the Coriolis effect, i.e., the deflection of a moving object in a rotating reference frame
- In the observer's view, the path of the moving object observed to be bent by a fictitious force







MEMS Gyroscope



SPRITZ
SECURITY & PRIVACY
RESEARCH GROUP



UNIVERSITÀ
DEGLI STUDI
DI PADOVA



AM10096V1



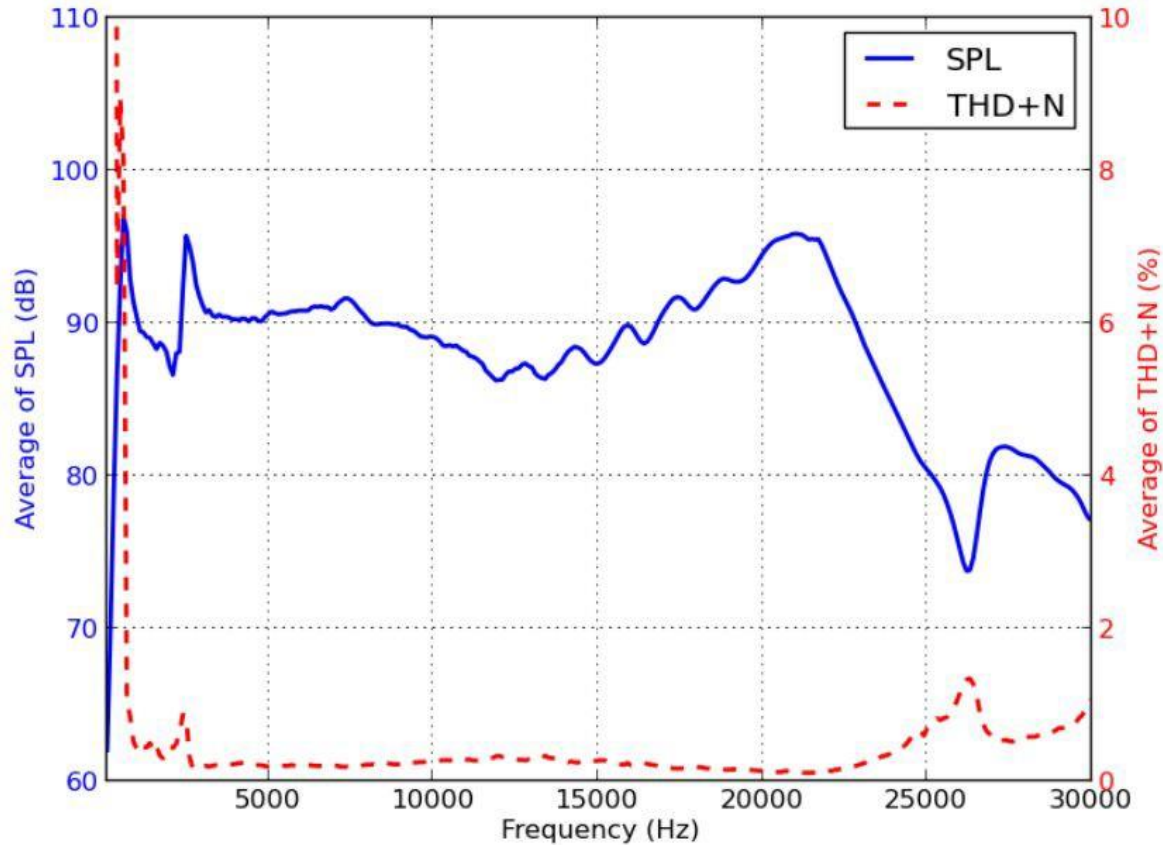
- MEMS gyroscopes are highly vulnerable to harsh acoustic noise
- Such noise cause accuracy degradation, thus providing wrong results
- A MEMS gyroscope has a resonant frequency related to the physical characteristics of its structure
- Due to this resonance, if met, the gyroscope generates unexpected outputs that cause system malfunctioning
- Usually, such frequencies should be ultrasound (above 20kHz)
- The sensitivity to noise can be exploited by an attacker



- To effectively attack MEMS with audio, we need their resonant frequencies
- A simple and reliable way, is exhaustive search → scan with single tone sound over a chosen frequency band
- Use a consumer grade speaker connected to a laptop and placed 10 cm above the top of the target gyroscope
- Generate single tone noises at frequencies from 100 Hz to 30 kHz at intervals of 100 Hz



- A common noise measurement unit for the loudness of sound is the Sound Pressure Level (SPL)
- Another important property is the Total Harmonic Distortion plus Noise (THD+N), i.e., the ratio of the power of the harmonics and noise components to that of a fundamental component (in %)
- Sound source = tweeter



- Gyroscopes fixed on a stable frame in an anechoic chamber
- In the absence of noise, the variance of the measurements should be zero → difference in standard deviation as a criterion for the resonance of the gyroscope

Sensor	Without noise			With noise			Ratio		
	$\sigma_{X_{wo}}$	$\sigma_{Y_{wo}}$	$\sigma_{Z_{wo}}$	σ_{X_w}	σ_{Y_w}	σ_{Z_w}	$\sigma_{X_w}/\sigma_{X_{wo}}$	$\sigma_{Y_w}/\sigma_{Y_{wo}}$	$\sigma_{Z_w}/\sigma_{Z_{wo}}$
L3G4200D	3.15	2.69	2.88	12.1	22.04	4.45	3.84	8.21	1.55
L3GD20	2.92	2.47	2.3	62.03	76.67	3.09	21.21	31.04	1.35
LSM330	13.09	16.03	21.45	177.71	114.34	30.44	13.57	7.13	1.42
MPU6000	11.79	13.92	12.8	12.48	14.74	111.21	1.06	1.06	8.69
MPU6050	13.21	12.32	11.17	13.8	12.55	58.17	1.04	1.02	5.21
MPU6500	17.34	19.63	18.21	363.21	71.04	56.15	20.95	3.62	3.08
MPU9150	10.69	11.47	10.71	10.98	11.97	58.59	1.03	1.04	5.47


Noise Effect on MEMS




SPRITZ
SECURITY & PRIVACY
RESEARCH GROUP



UNIVERSITÀ
DEGLI STUDI
DI PADOVA



L3GD20



MPU6000



- The attacker, in order to design a good strategy, needs to understand what the drone is doing
- Static analysis of the controller code to understand the reaction of sensor and actuators to noise
- Usually drones can support different gyroscope implementations, however the main software routing is the same for all of them

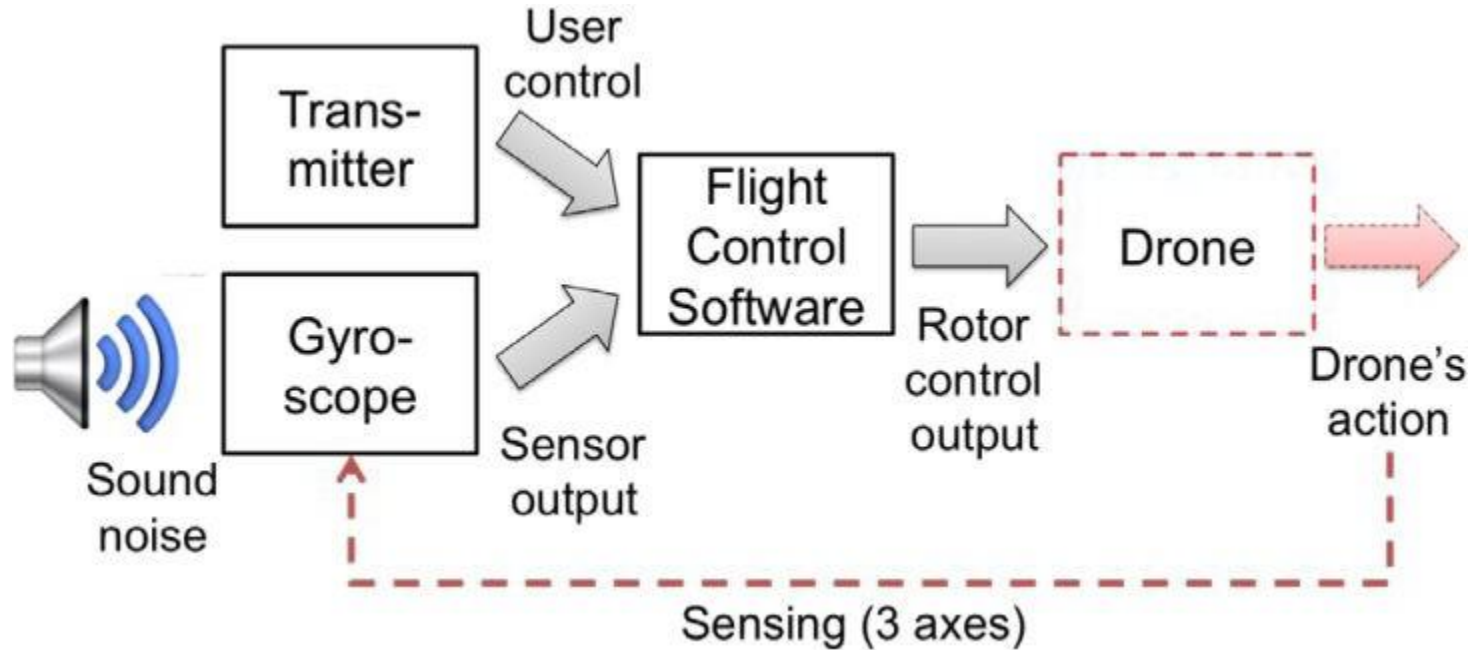


- The main processor reads the raw data from the gyroscope's registers through an I²C interface, along with the control data provided by the user
- Raw data for each axis is stored in two 8-bit register → main inputs of the controller, which uses a PID logic to instruct rotors as follows
 - P is proportional to the present output of the gyroscope, and if the present output is abnormally large the control from the transmitter can be ignored




- Such raw data are them main inputs of the controller, which uses a PID logic to instruct rotors as follows
 - I is proportional to the accumulated error between the output from the transmitter and the gyroscope, which can be ignored because usually its gain is very small
 - D is proportional to the changes between the present output values and the gyroscope
- Throughout the whole process the gyroscope data is not checked

Propagation of Signals

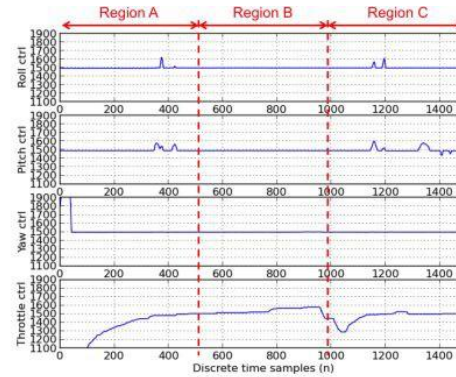




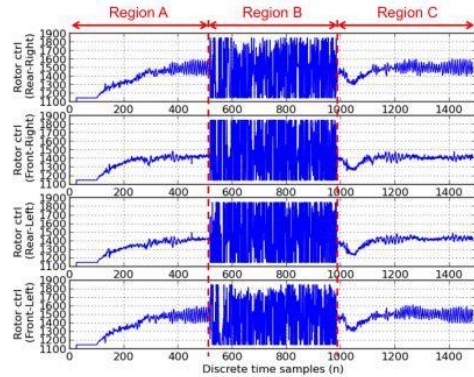
- Attach a small bluetooth speaker above the target system gyroscope at a distance of 10 cm to serve as attacking source
- The sound noise is turned on while the target drones were stably maintained in the air
- To observe the status of the drone before, during, and after the attack, the speaker is turned off, on, and off again every 10 seconds



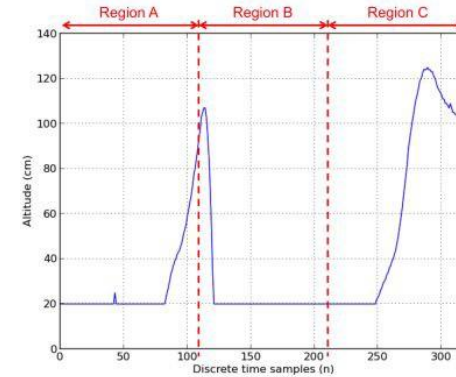
(a) Raw data samples of the gyroscope



(b) Received data samples from the transmitter



(c) Rotor control data samples (from the flight control software)



(d) Altitude data samples from sonar