UNICO I+D Project
6G-SORUS-DRONE

SORUS-DRONE-A3.1-E2 (E13)

Performance evaluation: final

Abstract

This deliverable advances the project's overarching goal of integrating UAVs, Reconfigurable Intelligent Surfaces (RIS), and virtualized RAN (vRAN) to enable intelligent, adaptive, and resource-efficient B5G/6G networks. By demonstrating how UAVs can autonomously map wireless coverage using probabilistic modeling and uncertainty-driven exploration, it establishes a critical foundation for mobility-aware network optimization. The generated coverage maps can inform RIS configurations, guide vRAN resource allocation, and support on-demand UAV deployments for dynamic connectivity extension. Together, these capabilities contribute to a coordinated orchestration framework where UAV-based sensing, RIS-enabled propagation control, and vRAN virtualization converge to deliver highly flexible, energy-efficient, and self-optimizing next-generation wireless networks.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia

Document properties

Document number	SORUS-DRONE-A3.1-E2
Document title	Performance evaluation: final
Document responsible	Pablo Serrano (UC3M)
Document editor	Pablo Serrano (UC3M)
Editorial team	Jonathan Almodóvar, Pablo Serrano, David Ramírez (UC3M)
Target dissemination level	Public
Status of the document	Final
Version	1.0
Delivery date	30-06-2025
Actual delivery date	30-06-2025

Production properties

Reviewers	Juan Manuel Montes, Marco Gramaglia (UC3M)

Disclaimer

This document has been produced in the context of the 6G-SORUS-DRONE Project. The research leading to these results has received funding from the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union-NextGenerationEU through the UNICO 5G I+D programme.

Contents

List of Figures

Executive Summary

This deliverable makes a significant contribution to the overall objective of integrating UAVs, Reconfigurable Intelligent Surfaces (RIS), and virtualized Radio Access Networks (vRAN) by focusing on the creation of high-fidelity wireless coverage maps. These maps, generated through UAV-based exploration combined with Gaussian Process modeling, provide a detailed, uncertainty-aware view of the radio environment, forming the foundation for data-driven orchestration in B5G/6G networks. By enabling UAVs to intelligently collect minimal yet highly informative measurements and progressively refine coverage maps, the proposed framework transforms UAVs into autonomous sensing agents capable of characterizing dynamic propagation conditions with precision.

Accurate coverage maps derived from this process are central to the project's broader vision. They can be leveraged to inform RIS configuration, allowing intelligent re-shaping of the propagation environment to address coverage gaps or enhance capacity in targeted regions. Similarly, these maps provide mobility-aware inputs to vRAN controllers, enabling more effective allocation of virtualized resources and adaptive cell reconfiguration based on real-time spatial demand. By linking UAV-based sensing with probabilistic modeling, this deliverable establishes a scalable and efficient mechanism to provide the situational awareness required for joint UAV, RIS, and vRAN coordination.

In doing so, this work bridges environmental sensing and network control, creating the analytical foundation for a closed-loop orchestration framework where UAV-derived coverage maps drive both RIS-assisted signal steering and vRAN optimization. This integration is critical for realizing self-optimizing B5G/6G networks capable of dynamically adapting to user distributions, traffic demands, and environmental conditions, ensuring resilient and energy-efficient connectivity in next-generation deployments.

Introduction

Recent advancements in aerial industry toward Unmanned Aerial Vehicles (UAVs) paved the way for a set of novel use cases in the sky, opening a new range of innovative applications. The varied sizes and shapes, coupled with the cost-effectiveness of UAVs, create opportunities in fields like package delivery, public safety, and medical support. UAVs can be broadly categorized into different types e.g., fixed-wing, rotary-wing, chopper drones, among others [SBM18].

The evolution of the aerial industry has seen a confluence with cutting-edge technologies like Reconfigurable Intelligent Surfaces (RIS), improving the capabilities of UAVs [LYL+22, YMZ+20, MMM+21]. RIS, with its ability to dynamically control and manipulate electromagnetic waves, plays a crucial role in enhancing communication and sensing capabilities of UAVs. By integrating RIS into the UAV systems, it becomes possible to optimize signal strength, mitigate interference, and adapt to dynamic environmental conditions, thereby significantly improving the overall performance and reliability of aerial operations. This synergy between UAVs and RIS not only extends the range of applications in fields such as surveillance and communication but also unlocks new possibilities in autonomous navigation and collaborative aerial missions [RTG+07].

This deliverable is organized as follows. Section 2 outlines the specific requirements guiding our analysis. These requirements encompass both technical specifications and contextual considerations, providing a clear framework for evaluating the effectiveness of UAVs and RIS in addressing the challenges posed by the evolving B5G landscape. Sections 3 and 4 delve into diverse scenarios and the subsequent performance evaluation, respectively. Section 3 is subdivided into three key subsections—UAV Radio, Control Capacity Planning, and Multiple UAVs—each delving into distinct aspects of the scenarios under consideration. These subsections lay the groundwork for our subsequent analysis by establishing the contextual framework. *UAV radio* delves into the intricate details of UAVs' radio communication aspects. By examining the UAV radio capabilities, we aim to identify strengths and limitations that play a pivotal role in shaping communication scenarios within B5G. *Control capacity planning* focuses on the packet-level dynamics, this subsection explores the efficiency of data transmission and reception in UAV networks. We investigate the packet-level intricacies to better understand the data flow and latency aspects crucial for B5G communication. *Multiple UAVs* envisions a multitude of interconnected devices, understanding the dynamics of multiple UAVs becomes imperative. This subsection explores the challenges and opportunities associated with the simultaneous operation of multiple UAVs, paving the way for robust communication strategies.

Section 4 evaluates the performance of UAVs and RIS in B5G scenarios. Subsections within this section explore critical KPIs, including Coverage Probability, Power Consumption, Control Dynamics and Data Operations, Cooperative UAVs, and RIS Gains. *Coverage Probability* delves into the coverage probability metrics, analyzing the extent to which UAVs and RIS contribute to ensuring reliable and expansive network coverage. Insights from this evaluation are essential for optimizing network

design to meet the demands of B5G scenarios. *Power Consumption* presents the power consumption patterns of UAVs and RIS, providing an understanding of their energy dynamics and paving the way for more energy-efficient solutions. *Control Dynamics and Data Operations* evaluates the reliability metrics, shedding light on the robustness and dependability of UAV and RIS-integrated networks. *Cooperative UAVs* investigates how cooperative efforts enhance overall system performance. This collaborative approach aligns with the collaborative nature of the envisioned B5G scenarios. *RIS Gains* focuses on the gains facilitated by Reconfigurable Intelligent Surfaces. By examining the impact of RIS on communication quality and efficiency, we aim to uncover the advantages these intelligent surfaces bring to the B5G landscape.

Through this deliverable, we highlight the intricate dynamics of UAV and RIS integration, providing a valuable resource for stakeholders navigating the evolving landscape of B5G communication networks.

1. Creating B5G/6G network coverage maps for better management

One particularly promising application is the use of UAVs for the generation of cartography maps that capture key coverage indicators such as Received Signal Strength Indicator (RSSI), Signal-to-Noise Ratio (SNR), and other essential metrics. This capability is especially critical in areas that are difficult to access by humans, as well as in scenarios where cost and time efficiency are paramount.

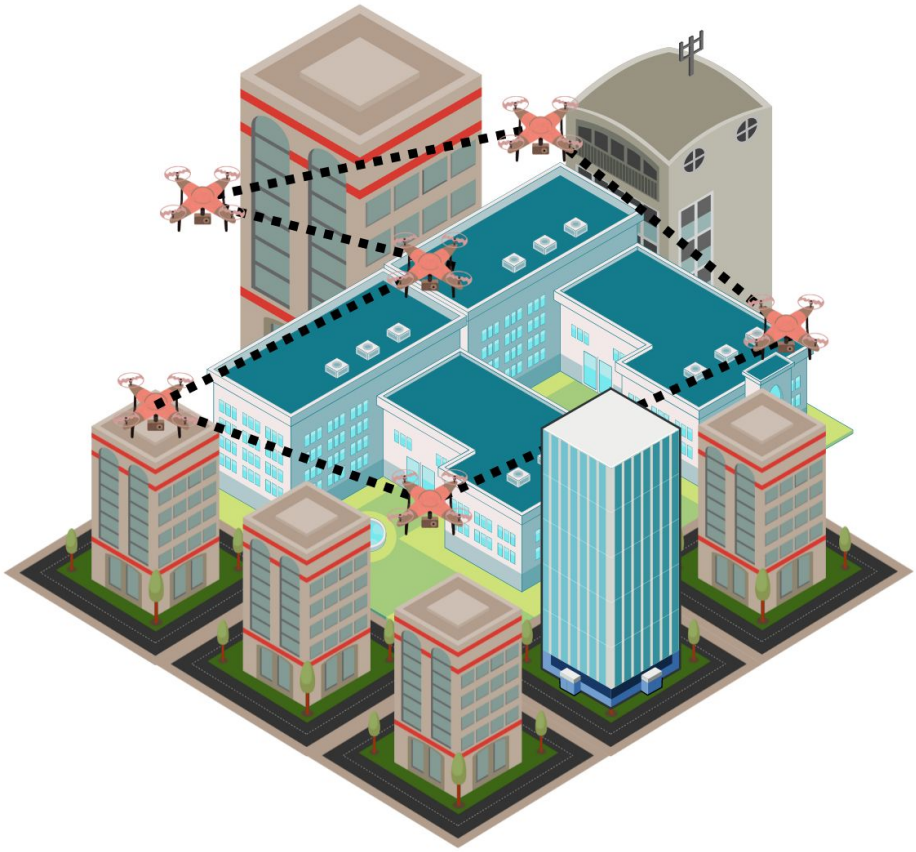

FIGURE 1: SCENARIO FOR UAV COVERAGE MAPS GENERATION.

Traditionally, coverage measurements and network optimization efforts rely on ground-based surveys that can be resource-intensive, time-consuming, and, in some cases, geographically limited. UAVs offer a transformative alternative, providing flexibility and precision while minimizing the challenges faced by human operators. Equipped with advanced sensors and communication technologies, UAVs can quickly and efficiently map large areas, gathering crucial data on network coverage in both urban and rural environments [1].

UAVs offer several key advantages in the field of coverage mapping, particularly in challenging environments. They excel in areas that are difficult or dangerous for humans to reach, such as mountainous terrains, industrial sites, or disaster zones, making them ideal for collecting network data where human deployment is impractical or unsafe. Deploying UAVs significantly reduces the resources required for traditional methods, as a single UAV can cover large areas in a fraction of the

time compared to ground-based teams, speeding up data collection while reducing operational costs. UAVs can also be pre-programmed to follow precise routes, ensuring full coverage with minimal human intervention. Their ability to fly at controlled altitudes and follow specific paths allows for high-precision measurements, capturing a wide range of coverage indicators like RSSI, SNR, latency, and throughput. UAVs' flexibility in hovering at specific locations and adjusting altitude enables granular data collection, resulting in highly detailed and accurate coverage maps essential for network optimization. In addition to static data, UAVs can provide real-time feedback, making them invaluable in dynamic environments. For instance, they can monitor network performance during events or in fast-changing conditions, such as natural disasters, allowing for immediate adjustments to network configurations and improving the overall resilience and adaptability of B5G networks [2,3].

2. Challenges

While the integration of UAVs for coverage mapping in B5G networks offers significant advantages, it also presents a range of challenges that need to be addressed for successful implementation. These challenges span technical, regulatory, and operational domains. From managing the complexities of real-time data transmission with ultra-low latency requirements to ensuring robust collision avoidance in increasingly congested airspaces, each aspect requires careful consideration. Moreover, challenges related to energy efficiency, path planning, and regulatory compliance add layers of complexity to UAV operations, particularly when deploying fleets of drones over large areas or in dynamic environments. In the following sections, we will delve into these challenges, exploring the technological and strategic hurdles that must be overcome to fully harness the potential of UAVs in B5G networks, [4].

2.1. Collision avoidance and latency

Latency and collision avoidance are two critical aspects in the deployment of UAVs for coverage mapping in B5G networks, both of which have been explored in depth in previous deliverables. When it comes to latency, the ability of UAVs to capture and transmit real-time data is essential for network operators who rely on immediate feedback to adjust configurations or identify coverage issues. In B5G networks, ultra-low latency is a key requirement, ensuring that UAVs can provide instantaneous reporting of coverage indicators such as RSSI or SNR without delay. This is particularly important in fast-changing environments, such as during disaster recovery or large-scale events, where network conditions can fluctuate rapidly, [5, 6].

However, ensuring low latency is only part of the challenge. Collision avoidance is equally crucial, particularly when deploying multiple UAVs in the same airspace. The risk of collisions not only endangers the UAVs themselves but also compromises the quality and reliability of the data they are collecting. Effective collision avoidance mechanisms are vital for maintaining continuous operation, particularly in complex or congested environments such as urban areas or emergency situations. Previous deliverables have discussed various strategies for collision avoidance, including the use of onboard sensors, machine learning algorithms, and communication protocols that allow UAVs to share positional data and adjust their trajectories in real time. These technologies ensure that UAVs can operate safely near one another, avoiding disruptions to data collection and preserving the integrity of the network mapping process [7,8].

2.2. Path planning

Path planning and policy management are integral to the successful deployment of UAVs in B5G networks, particularly when it comes to ensuring efficient coverage mapping and data collection. In previous deliverables, specifically in the state-of-the-art analysis of Deliverable X, we explored various algorithms designed to optimize UAV trajectory and task execution. Among these, Deep

Reinforcement Learning (DRL) has emerged as one of the most promising techniques for UAV path planning, offering the adaptability and real-time decision-making capabilities required for complex environments [9,10].

DRL-based algorithms excel in dynamically adjusting UAV flight paths based on real-time conditions, such as changes in network coverage, obstacles, and airspace restrictions. By leveraging the continuous learning framework of DRL, UAVs can autonomously refine their trajectories to maximize data collection efficiency, avoid obstacles, and reduce power consumption. This is especially valuable in scenarios where coverage mapping must be done across large or complex terrains, where traditional pre-defined flight paths may not be optimal. DRL allows UAVs to navigate such environments more effectively, adjusting to unforeseen circumstances and learning from previous experiences to improve performance over time.

In policy management, DRL enables the UAV to make intelligent decisions based on multiple factors, such as minimizing latency, ensuring collision avoidance, and optimizing energy consumption. For example, the UAV can prioritize specific areas for coverage based on real-time data, such as regions where network signal strength is weak or where users are experiencing connectivity issues. Additionally, DRL-based systems can enforce policies related to altitude control, speed, and data transmission, ensuring that UAVs operate within predefined parameters while still adapting to on-the-ground realities [11].

One of the strengths of DRL in UAV path planning is its ability to consider long-term rewards rather than immediate gains. This is particularly useful in scenarios where the UAV needs to balance multiple competing objectives, such as maintaining coverage, avoiding obstacles, and conserving energy. The ability of DRL to weigh these factors and adapt its decisions dynamically allows for more efficient and reliable UAV operations, making it a valuable tool for optimizing network coverage in B5G environments.

3. Problem Definition

Having covered the technical aspects of UAV deployment in earlier deliverables, we'll now focus on the data collection process, particularly how we can develop an efficient framework for gathering key coverage indicators like RSSI (Received Signal Strength Indicator). The way we model the environment is essential in guiding UAVs to collect data efficiently, especially when they're mapping network coverage over complex or remote areas.

The idea is to create a system that allows UAVs to gather RSSI data while considering factors like terrain, obstacles, and how signals behave in different environments. By having a detailed model of the area, UAVs can be directed to zones where the signal is likely to be weaker or inconsistent, focusing on these problem areas to get a clearer picture of network performance. This makes the coverage maps more accurate and helps avoid unnecessary measurements, saving both time and battery life.

In practice, this means UAVs will be able to adjust their paths as they collect data, zeroing in on areas with low signal or where they detect sudden changes in strength. The environment model will factor in things like buildings, trees, or other physical barriers that might affect signal propagation, ensuring that the UAVs can navigate around them and still collect useful data. By focusing on smarter, more targeted data collection, this approach ensures that UAVs can gather the most valuable information with minimal effort, ultimately helping improve the coverage and quality of B5G networks [12], and recent efforts such as UAV experimentation within O-RAN testbeds further demonstrate the feasibility of integrating UAV-based measurements into operational RANs [13].

3.1. Measurements model

To model the signal strength measurements, such as RSSI, along a UAV's path, we will employ **Gaussian Process (GP) models**. These models are highly effective for predicting spatially continuous phenomena because they can provide both the predicted value and the uncertainty at each location. This makes them particularly useful in creating smooth coverage maps based on discrete UAV measurements, filling in gaps where direct measurements are unavailable [14].

Gaussian Processes in One Dimension

A Gaussian process is a collection of random variables, any finite number of which have a joint Gaussian distribution. In simpler terms, it's a generalization of the Gaussian (normal) distribution, not just for random variables but for functions. Mathematically, we define a Gaussian process as:

$$f(x) \sim GP\big(\mu(x), k(x, x')\big)$$

Where:

- $f(x)$ is the function we are trying to model (e.g., RSSI as a function of location),
- $\mu(x)$ is the mean function, which represents the expected value of $f(x)$ at any point x,

- $k(x, x')$ is the covariance function (or kernel), which determines how closely related two points x and x' are.

In a one-dimensional Gaussian process, x represents the position along a path (for example, a UAV's flight path), and we assume that the values of $f(x)$ at nearby points are more correlated than those far apart. This is captured by the kernel function $k(x, x')$, which measures the covariance between the function values at two points x and x'. A common choice for the kernel is the squared exponential (or radial basis function) kernel:

$$k(x, x') = \sigma_f^2 \exp\left(-\frac{(x - x')^2}{2l^2}\right)$$

Where:

- σ_f^2 is the variance (determining the overall scale of the function's variation),
- l is the length scale (which controls how rapidly the function can change).

Given a set of observations $\{x1, y1\}, \{x2, y2\}, \ldots, \{xn, yn\}$, where $y_i = f(x_i) + \epsilon y_i = f(x_i) + \epsilon$ $(with\ \epsilon \sim N(0, \sigma_n^2))$ being measurement noise), the GP model allows us to predict the value of $f(x)$ at a new location x^* using the following properties:

Mean of the Prediction:

$$\mu^*(x^*) = \mu(x^*) + k(x^*, X)[K(X, X) + \sigma_n^2 I]^{-1}\big(Y - \mu(X)\big), \qquad (1)$$

Where:

- X is the vector of input locations,
- Y is the vector of observed function values,
- $k(x^*, X)$ is the covariance between the new point x^* and the observed points,
- $K(X, X)$ is the covariance matrix between the observed points.

Variance of the Prediction:

$$\sigma^{*2}(x^*) = k(x^*, x^*) - k(x^*, X)[K(X, X) + \sigma_n^2 I]^{-1} k(X, x^*), \quad (2)$$

This allows us to not only predict the value of the function at any new location x^* but also **to quantify how uncertain the prediction is based on the available data**. The covariance function $k(x, x')$ ensures that the predicted values are more strongly influenced by nearby data points.

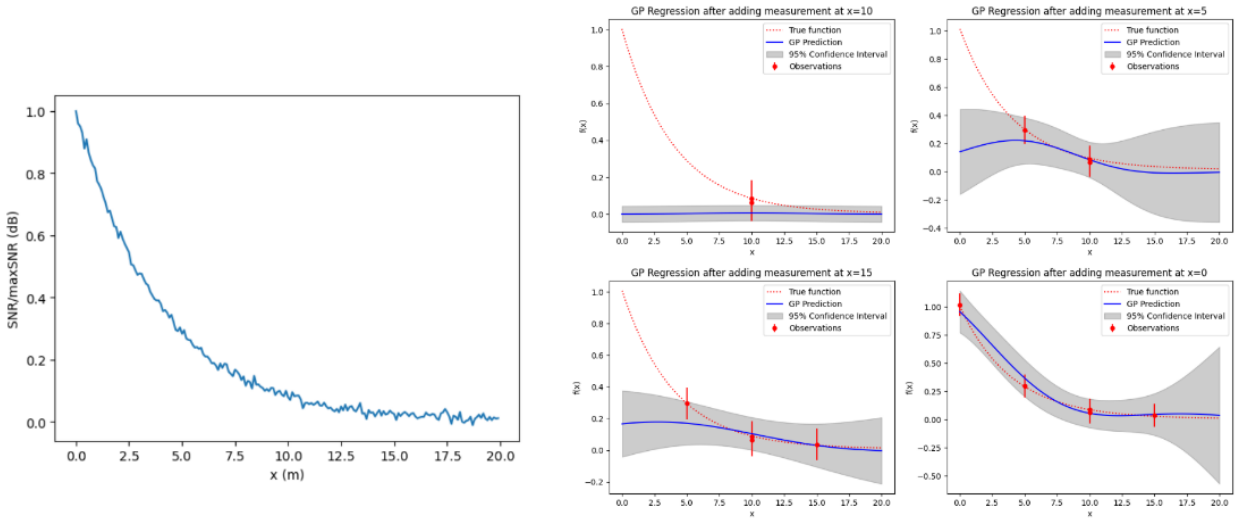

FIGURE 2: TOY EXAMPLE OF GP MEASUREMENT PROCESS IN 1D

Figure 2 illustrates the application of the Gaussian Process (GP) framework described in the previous section to model signal strength (e.g., RSSI or SNR) along a one-dimensional UAV path. The left plot shows the true normalized signal decay with distance. The four subplots on the right demonstrate how the GP posterior is updated as new measurements are added sequentially at $x = 0, 5, 10$, and 15 meters.

Each update follows the predictive mean and variance equations introduced earlier (Equations (1) and (2)). Initially, with no data, the model relies solely on the prior, leading to high uncertainty. As new observations are incorporated, the posterior mean becomes more accurate and the confidence intervals shrink, particularly near the observed points—reflecting the spatial correlation encoded by the kernel function $k(x, x')$. This exemplifies how GPs can adaptively refine predictions and quantify uncertainty in UAV-based signal mapping.

Gaussian Processes in Two Dimensions

Extending to a 2D Scenario for Realistic UAV Data Collection

While a one-dimensional Gaussian process provides valuable insights along a single UAV flight path, a more realistic scenario for coverage mapping involves a two-dimensional (2D) environment. In this case, we consider a grid where each measurement taken by the UAV at a given location not only provides information about that specific point but also affects the predictions at nearby, contiguous points. This approach models real-world conditions more accurately, as signals like RSSI typically exhibit spatial correlation across both dimensions, latitude and longitude in geographic space [15].

Gaussian Processes in Two Dimensions

When we shift to a more realistic 2D scenario for UAV data collection, things get a bit more practical. Instead of just moving along a single path, UAVs cover wider areas, flying over a grid of locations.

The goal here is to understand how signal measurements taken at one point can help predict what's happening in the surrounding areas, giving us a better overall map of network coverage.

In this setup, each time the UAV takes a measurement, that data doesn't just tell us about the signal strength at that exact spot—it also gives clues about what's going on nearby. For instance, if the signal is strong at one location, it's likely that the signal will be somewhat strong in the surrounding areas too, but it may taper off the farther out you go. By using this kind of logic, we can predict coverage across the grid, even in places where no direct measurements were taken.

This approach allows us to gather enough data without having to measure every single point. The model helps fill in the blanks by considering how the signal behaves across space, making sure that the points close together influence each other. This gives us a more complete view of the area's coverage, which helps with planning and optimizing network performance.

The Uncertainty Matrix

For each node in the 2D grid, in addition to the predicted RSSI value, we also have an uncertainty value. This is stored in the form of an uncertainty matrix, where each entry corresponds to a grid node and reflects the variance $\sigma^2(x)$ at that specific point. The uncertainty matrix provides a clear visual representation of which areas of the grid are well understood (low uncertainty) and which areas require more data collection (high uncertainty).

This uncertainty matrix is vital because it will help us develop the positions the UAV must prioritize. For example, in regions where the uncertainty is high, the UAV might need to take more measurements to reduce this uncertainty and improve the accuracy of the coverage map. Conversely, in regions where the uncertainty is low, fewer additional measurements are necessary.

4. Methodology

Figure 3 illustrates the overall workflow of a UAV-based system designed to efficiently map network metrics across a two-dimensional area. The goal of this system is to intelligently explore the environment, collect a minimal yet informative set of measurements, and use a probabilistic model to infer a complete coverage map. This process is structured as a closed-loop interaction between the UAV, the environment, and the learning algorithm, enabling adaptive and efficient data collection.

The process begins with the environment, depicted in the leftmost part of the diagram. The environment contains a spatial distribution of network signal strength, which is initially unknown to the UAV. In the visualization, this environment is represented by a heatmap on a grid, where each cell corresponds to a location in the two-dimensional area. Signal sources, such as base stations or access points create overlapping regions of coverage, generating a non-uniform RSSI landscape across the grid. This underlying signal distribution is what the UAV aims to discover and model.

To achieve this, the UAV performs a series of measurements at selected grid points. These measurements are collected as the UAV flies over the area and are represented by red "X" markers on the second heatmap. Each marker indicates a location where the UAV has recorded an RSSI value. These discrete observations are passed to a probabilistic model. In this case, a Gaussian Process (GP) which is responsible for learning the spatial relationships between the measured points and interpolating the signal strength in unvisited regions. The GP model leverages the assumption that physically close locations are likely to exhibit similar signal characteristics, allowing it to produce a smooth and continuous estimate of the entire signal field, even in areas without direct measurements [16].

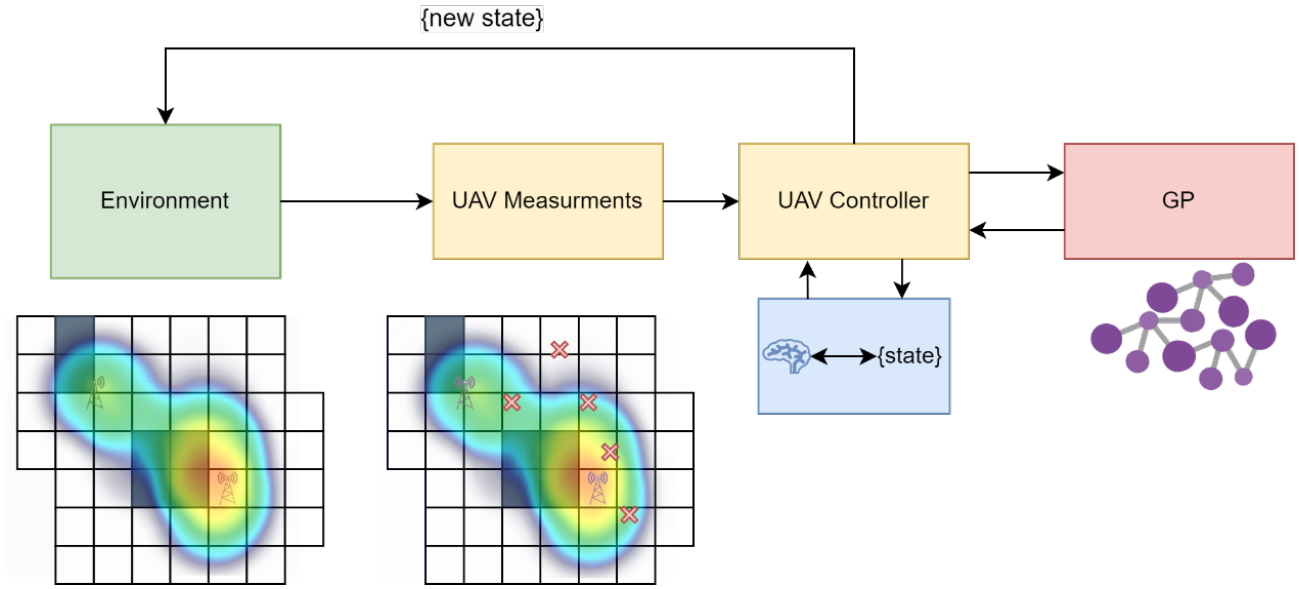

FIGURE 3: WORKFLOW OF THE COVERAGE MAP PROCESS

The predictions generated by the GP are then used to define the current system state, which includes both the estimated coverage map and the associated uncertainty in different regions. This state information is sent to a UAV controller module, which determines the next action the UAV should take. The controller may use an exploration strategy such as uncertainty-based sampling or reinforcement learning to select the most informative next location, balancing the need to explore unknown areas with the goal of refining the overall map. The chosen action corresponds to a new measurement location for the UAV, which it then visits to obtain another RSSI reading.

Once the new measurement is acquired, it is added to the dataset, and the GP model is updated with the latest information. This update yields a new state, incorporating the refined predictions and reduced uncertainty. The loop then repeats: the updated state is passed to the controller, which selects the next action, and the UAV continues its exploration. This iterative process allows the UAV to progressively build a high-fidelity signal coverage map while minimizing the number of physical measurements required.

The figure also emphasizes that while the example operates over a structured 2D grid, the approach can generalize to more complex domains modeled as graphs. Each node in such a graph may represent a measurement location such as a street intersection or a critical waypoint in rugged terrain while edges capture the spatial or semantic relationship between locations.

Overall, this closed-loop system demonstrates how UAVs can be used for efficient and intelligent data collection in wireless networks. By combining sparse measurements, probabilistic modeling, and adaptive control, it is possible to construct detailed coverage maps that support a wide range of applications, from network planning to real-time monitoring.

4.1. UAV Coverage Mapping

For UAVs collecting data over a 2D area, the goal is to build a comprehensive map of RSSI (or any other network metric) across the region. The UAV flies over the area, collecting data at discrete locations, which are treated as grid nodes. The GP model then interpolates the RSSI values at all other grid points, using the spatial correlations between nearby locations to predict signal strength in unmeasured areas.

This approach allows for efficient data collection since the UAV doesn't need to measure every single point on the grid. Instead, it can collect data at strategic locations, and the GP will fill in the gaps based on the correlation between the grid points. As a result, we can generate a highly accurate and continuous map of network coverage with fewer measurements, saving both time and resources while maintaining high-quality predictions across the 2D space.

The environment we've described as a simple 2D grid is just one example of a graph structure, where each node (representing a measurement point) is connected to its neighboring nodes. In more complex scenarios, this concept can be generalized to arbitrary graphs, allowing for greater flexibility in modeling UAV-based data collection in real-world environments. Unlike a regular grid, nodes in a general graph can represent various points of interest, and edges between them can vary in distance,

direction, or significance, depending on the landscape. For example, nodes could be street intersections in an urban area or vantage points in mountainous terrain, with edges weighted by factors like distance or difficulty. The Gaussian process framework can adapt to these complex graphs by using graph distances instead of simple physical distances to model relationships between points, enabling more accurate signal propagation predictions in real-world settings.

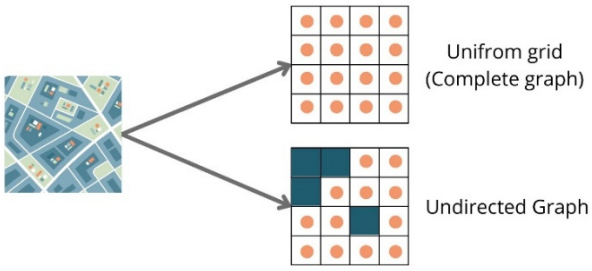

FIGURE 4: GRAPH CONSTRUCTION

Optimizing Measurement Locations on a Graph

By treating the environment as a graph rather than a strict grid, we can strategically choose where to place measurement nodes to better fit the setting. In practice, this means that UAVs can be programmed to take measurements at key nodes on the graph, rather than blindly following a grid pattern. For example:

In densely populated areas, measurements might be concentrated at important intersections where network performance is critical.

In sparsely populated regions, fewer nodes may be required, but these nodes could be positioned to cover key points like high-traffic roads or remote access points.

By designing a custom graph that reflects the actual environment and network needs, we can optimize the UAV's flight path and data collection process. The Gaussian process will still interpolate the signal strengths across the graph, predicting values at nodes where no measurements were taken, but now the model fits the specific characteristics of the area.

4.2. Collection process models

In our approach to collecting RSSI data for network coverage mapping, the UAV is assumed to operate with two key actions: it can either move to a nearby node in the grid or stay at its current location to take more precise measurements. This dual strategy of balancing between movement and focused measurement, allows for flexible and adaptive data collection, improving both the quality of individual measurements and the efficiency of overall coverage.

The UAV moves across a predefined graph, with each node representing a location where RSSI measurements can be collected. Movement between nodes is assumed to occur at a constant speed,

which simplifies the modelling of flight dynamics. This speed ensures that the UAV can efficiently cover large areas of the grid, making it suitable for wide-area coverage mapping.

By moving from node to node, the UAV can quickly gather a general sense of the network's signal strength over the entire region. This is particularly useful in areas where signal propagation is relatively uniform, as fewer measurements at each node may be sufficient to construct an accurate map. Movement-based data collection allows for faster coverage of the grid and helps to identify general trends in network performance across large areas.

In some cases, rather than moving to a new location, the UAV might benefit from staying at its current node for a longer period to take additional measurements. This is particularly valuable in areas where signal strength is highly variable or where the initial measurements show significant uncertainty.

By staying at the current node, the UAV can collect more samples of the RSSI, improving the accuracy of the measurement at that specific point. This can reduce the impact of temporary fluctuations in signal strength, providing a more reliable and stable estimate of the RSSI at the given node. More precise measurements also help to reduce the uncertainty associated with that node, which is reflected in the uncertainty matrix discussed earlier.

Heuristic algorithm

To begin the measurement process on the graph, we first design a heuristic algorithm that focuses on moving the UAV to nearby nodes where the uncertainty of the signal strength is high. The core idea is simple: as the UAV collects data, the Gaussian Process model estimates the signal strength across the grid, but the predictions for nodes without direct measurements will naturally have higher uncertainty.

Our heuristic algorithm leverages this by prioritizing the nodes with the highest uncertainty that are closest to the current position of the UAV. After taking a measurement at one node, the UAV then moves to a neighbouring node where the model predicts the greatest uncertainty, refining the coverage map as it goes. This approach helps ensure that the UAV efficiently reduces uncertainty in local areas, incrementally improving the accuracy of the overall map without having to cover large distances between measurements.

The model we've developed has two versions, each with a different level of decision-making ability for the UAV, depending on how much of the surrounding area it can assess.

In the first version, the UAV can only "see" the uncertainty levels of its immediate neighbours. After it takes a measurement, it looks at the nearby nodes and moves to the one with the highest uncertainty. This keeps the UAV's decision-making simple and focused on what's right in front of it, allowing it to make quick, localized decisions.

The second version gives the UAV a broader view by allowing it to consider not just the immediate neighbours, but also the "neighbours of neighbours." This wider perspective helps the UAV make more informed choices, even if it means flying a little further. It can see where uncertainty is highest

in a larger area and can prioritize those spots, which might speed up the process of reducing overall uncertainty across the network.

Reinforcement Learning

Deep reinforcement learning lifts the coverage-mapping task from a rule-driven routine to an experience-driven optimisation problem, enabling the UAV to discover sophisticated behaviours that a hand-crafted heuristic is unlikely to capture. Whereas a static rule set must strike a single, conservative balance between "measure" and "move" for every terrain, propagation environment and battery budget, a DRL policy watches the evolving Gaussian-process uncertainty landscape in real time and adjusts its strategy step-by-step, choosing to linger only where the expected information gain per joule is high and otherwise pressing forward to unexplored regions. Because the reward explicitly prices both variance reduction and resource expenditure, the agent learns subtle trade-offs, such as skipping low-value detours, clustering extra samples along fast-fading canyon walls, or exploiting tailwinds to reach distant high-uncertainty pockets before the battery sags, that are invisible to a fixed threshold. Over dozens of simulated flights, the neural policy internalises the spatial statistics of path loss, the noise floor of its own receiver and even the cost of sharp manoeuvres, so that when it is deployed in a new but similar environment it immediately exhibits near-optimal behaviour without manual retuning. Moreover, the function-approximation capacity of DRL lets the agent extrapolate beyond the training graph: if sudden interference spikes inflate variance in a previously "mapped" sector, the policy can recognise the anomaly, return for a quick verification sweep, and then resume its mission, all without any explicit contingency code. In short, DRL converts the UAV from a deterministic sampler into an adaptive explorer that continuously weighs information value against operational cost, delivering denser maps where they matter, leaner traversals where they do not, and uniformly higher mapping quality per unit flight time.

A **deep-reinforcement-learning (DRL) formulation** treats the coverage-measurement mission as a finite-horizon Markov-decision process. The **agent** is the UAV itself; every decision step t it perceives a **state vector**

$$s_t = [v_t, \mu_t(V), \sigma_t(V), e_t, t]$$

where v_t is the index of the current graph node, $\mu_t(V)$ and $\sigma_t(V)$ are, respectively, the GP posterior mean and variance (or compact summaries such as "maximum remaining variance" and a binary mask of "mapped" nodes), and e_t is the remaining energy or flight time. Because the world is naturally a graph $G = (V, E)$, we embed s_t with a graph neural encoder so that spatial correlations and neighbourhood structure are preserved before feeding the result to the policy/value networks.

The **action set** is deliberately small to mirror the heuristic:

1. **Measure** – remain at v_t and acquire a new RSSI sample $y_t = f(v_t) + \epsilon$, updating the GP posterior with Equations (1)–(2).
2. **Move** – choose one adjacent node $v' \in \mathcal{N}(v_t)$ and fly there, consuming time and energy but adding no new data.

Transitions are straightforward: a *measure* action leaves the position unchanged and updates the posterior; a *move* action updates the position and decrements the energy budget while the posterior remains the same.

The **reward** encourages rapid, informative mapping while discouraging wasteful motion:

$$r_t = [I_{tot}(t) - I_{tot}(t+1)] - \lambda_m 1_{\{move\}} - \lambda_e \Delta e_t$$

where $I_{tot}(t) = \sum_{v \in V} \sigma_t^2(v)$ is the total posterior uncertainty, the first term therefore measures *information gain*, and λ_m, $\lambda_e > 0$ weight the costs of motion and energy usage.

The **objective** is to maximise the expected discounted return:

$$\max_{\pi_\theta} \mathbb{E}_{\pi_\theta} \left[\sum_{t=0}^{T} \gamma^t r_t \right]$$

with discount factor $\gamma \in (0,1]$. Because the action space is discrete and small, algorithms such as Deep Q-Networks are adequate, but actor-critic methods (e.g. PPO) offer smoother updates and, paired with graph encoders, scale to larger maps. The policy network $\pi_\theta(a \mid s)$ learns when the marginal value of an extra measurement outweighs its opportunity cost, and when it is more profitable to explore high-uncertainty regions.

Episodes terminate either when the map meets the confidence goal or when the time-/energy budget is exhausted. Once training converges, the learned policy generalises to unseen areas with similar propagation statistics, adapting measurement density on-the-fly instead of relying on fixed heuristic thresholds.

4.3. Coverage map creation

The performance of the SorusBoxScan platform was validated through a real-world measurement campaign designed to assess its ability to produce accurate coverage maps. For this purpose, the device was mounted inside a vehicle and operated alongside a commercial smartphone, the Samsung Galaxy S23 Ultra, running the G-NetTrack Lite application. Running both systems in parallel under identical conditions provided a robust methodology for cross-validation. By comparing the measurements obtained with the SorusBoxScan against those from a widely used industry tool, the experiment ensured that the platform's results could be trusted for reliable coverage analysis.

During the campaign, the vehicle followed a predefined route that traversed areas with different coverage conditions, including strong outdoor signals, weaker regions near urban obstacles, and sections where mobility triggered handovers between cells. The SorusBoxScan collected metrics through its integrated workflow, publishing them via MQTT, while the smartphone logged its readings directly through the application. After the campaign, both datasets were processed to generate coverage maps showing the average Reference Signal Received Power (RSRP) values along the route. A shared color scale was applied, ranging from red for weak coverage (below –100 dBm) to green for strong coverage (better than –80 dBm), enabling a direct visual comparison between the two devices.

The resulting maps reveal a high level of consistency between the smartphone and the SorusBoxScan platform. In both cases, strong coverage areas appear in green, while weaker zones, such as those around Parla and near the tunnel sections, are clearly visible in orange and red. This alignment demonstrates that the SorusBoxScan system can replicate the behavior of a commercial smartphone and, crucially, transforming its measurements into coverage maps that faithfully represent real-world network conditions. A closer look also shows that the SorusBoxScan maps present smoother transitions between coverage levels, suggesting that its data handling pipeline produces stable results even under rapidly changing conditions.

An additional advantage of the methodology was the ability to collect measurements across multiple frequency bands, including LTE bands b0, b3, b7, b8, and b20, as well as the 5G New Radio band n78. When mapped, these results highlight the distinct characteristics of each band: low-frequency bands like b8 and b20 offered wide coverage and strong penetration, while higher-frequency bands such as b7 and n78 provided higher capacity but with a more fragmented footprint. By representing these differences visually, the maps clearly illustrate the trade-off between coverage and capacity across technologies. This multi-band mapping capability makes the SorusBoxScan platform a versatile tool for evaluating modern network deployments.

The creation of these coverage maps also demonstrates the flexibility of the data pipeline. Measurements can be monitored live using an MQTT client, enabling immediate visualization of coverage during a campaign. They can also be processed programmatically through the SorusBoxScan SDK, making it possible to automate map generation or integrate the results into custom analysis workflows. Finally, by storing results in a database, maps from different campaigns can be aggregated to build long-term coverage models and track how performance evolves over time.

In summary, the campaign confirmed that the SorusBoxScan platform can reliably generate coverage maps that match those obtained with a commercial smartphone, while also offering greater flexibility for data processing and integration. Its ability to capture multi-band behavior, operate robustly under mobility conditions, and provide reusable data workflows makes it a practical and powerful solution for researchers and operators interested in visualizing and optimizing mobile network coverage.

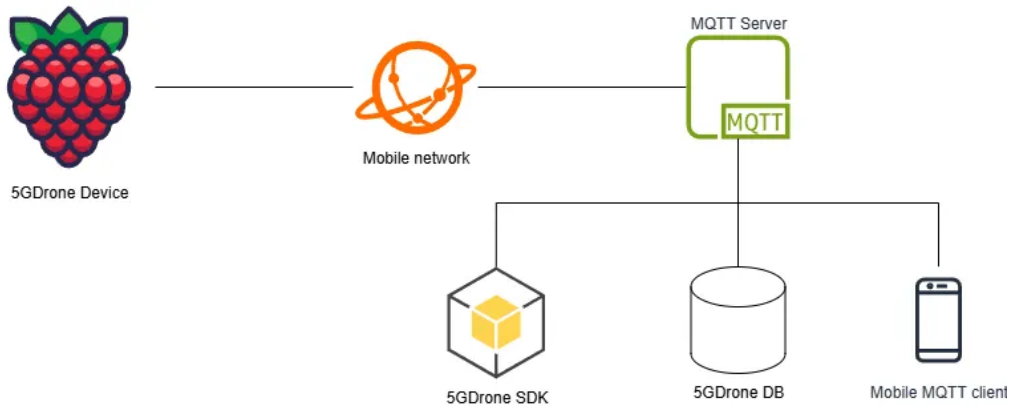

FIGURE 5: SORUSBOXSCAN DATA COLLECTION AND PROCESSING WORKFLOW

5. Evaluation

5.1. Coverage map creation - Gaussian Process

In this section, we present the results of our evaluation, focusing on the heuristic model among the two proposed collection process strategies. While deep reinforcement learning (DRL) offers clear advantages in terms of adaptability and experience-driven optimisation, its evaluation requires extensive training episodes, significant computational resources, and precise environment modelling to ensure reliable convergence and reproducibility. Given the scope of this deliverable, which prioritises validating the feasibility of our data collection framework and providing actionable insights under realistic constraints, we opted to assess the heuristic model first. This choice allows us to establish a strong performance baseline using a lightweight, interpretable strategy that does not depend on specialised training infrastructure. Moreover, by isolating the heuristic model, we can clearly quantify the improvements that DRL-based methods may bring in future work, where their ability to dynamically balance exploration and energy expenditure can be rigorously benchmarked against this baseline. The following results therefore focus on the heuristic approach, providing a foundation for subsequent comparisons with DRL-enhanced policies.

The environment used for evaluating the collection process models is illustrated in Figure 5. It is represented as a 2D grid where each cell corresponds to a discrete location in the area of interest, and the color intensity reflects the underlying signal strength or measurement value (e.g., received power or information density) at that point. This environment includes two distinct high-value regions (visible as bright yellow-green clusters) embedded within a lower-value background. These regions simulate areas of higher RSSI where additional measurements are particularly valuable.

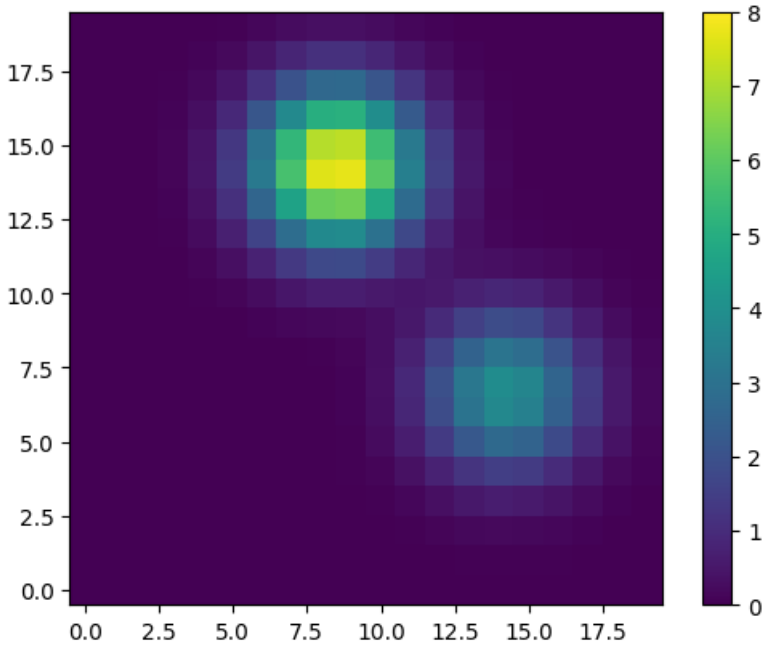

FIGURE 6: SIMULATION ENVIRONMENT

The first algorithm follows a simple exhaustive approach: the UAV sequentially visits every grid node in the environment and collects a measurement at each location. This method provides a comprehensive coverage of the area and establishes a performance baseline for evaluating more advanced collection strategies.

Figures 6 and 7 illustrate the evolution of the predicted Signal-to-Noise Ratio (SNR) and its associated uncertainty across the environment as measurements are progressively gathered.

- **Initial state:** At the beginning (Figure 6), no measurements have been taken. The predicted SNR is uniformly uninformative, and the uncertainty is maximal across the entire grid, reflecting the absence of prior knowledge.
- **Intermediate stage:** As the UAV moves and collects measurements (Figure 7), the predicted SNR begins to align with the underlying environment in the explored regions. Uncertainty decreases locally around visited nodes but remains high in unvisited areas.
- **Final state:** After completing the full traversal (Figure 8), the predicted SNR closely matches the true environment, and uncertainty is minimized throughout the grid. This demonstrates that exhaustive sampling effectively maps the environment but incurs significant flight time and energy cost.

While this baseline guarantees maximum accuracy, its inefficiency underscores the need for smarter strategies that selectively sample high-value regions or dynamically adapt flight paths to reduce resource expenditure while preserving mapping fidelity.

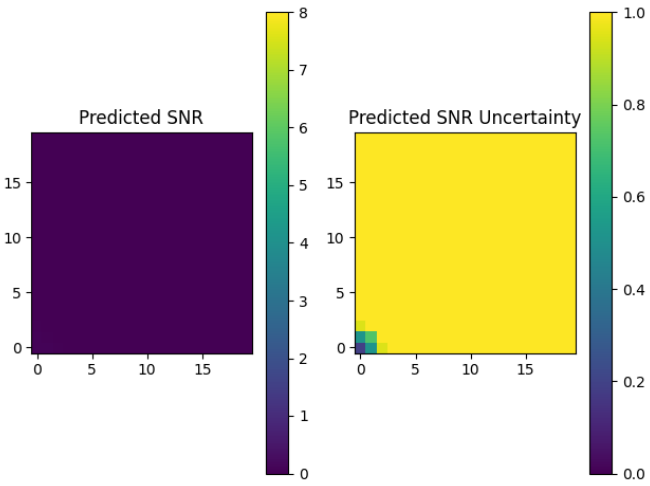

FIGURE 7: INITIAL STATE

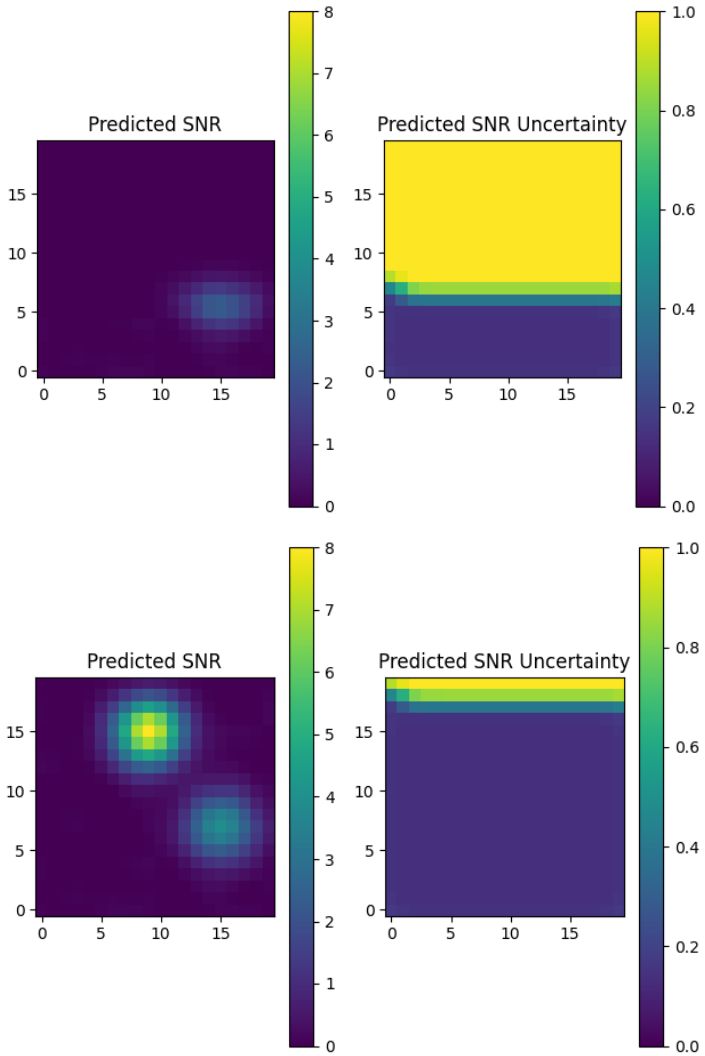

FIGURE 8: RESULTS OF THE PASSING THOUGH ALL THE NODES

Figures 6-8 illustrate the progression of the heuristic algorithm, which selectively samples nodes based on uncertainty thresholds adjusted for their position (corner, edge, or inner) and evaluates both immediate and second-neighbor uncertainties when planning its next move.

- **Initial state:** The same as the one inf Figure 6.
- **Intermediate stage:** In Figure Y, the UAV begins to traverse the grid more strategically, focusing on areas where uncertainty remains concentrated. The predicted SNR gradually aligns with the true underlying map, particularly around the two high-value regions, while large swathes of low-value areas remain unsampled. Importantly, the second-neighbor evaluation enables more efficient planning: rather than exhaustively sampling every node, the UAV skips low-information points and moves toward regions where both direct and second-neighbor uncertainties indicate high potential for variance reduction.
- **Final state:** By the end of the run (Figure Z), the predicted SNR closely matches the environment while the uncertainty has dropped significantly, especially around the two

Gaussian peaks. Although some residual uncertainty remains in less critical zones, this approach achieves comparable mapping fidelity with substantially fewer measurements and reduced traversal length compared to the exhaustive baseline.

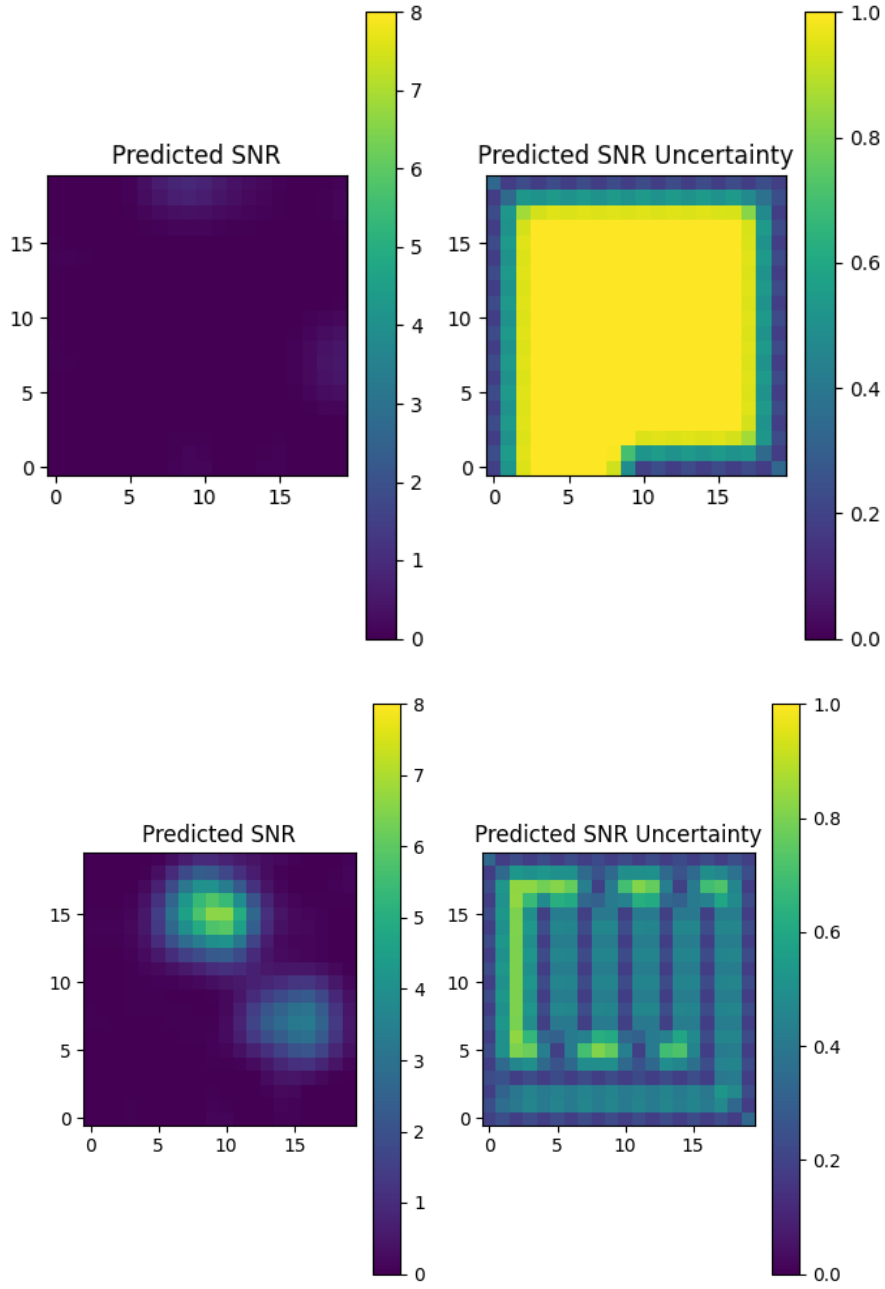

FIGURE 9: RESULTS OF THE ALGORITHM KNOWING THE UNCERTAINTY

Overall, the heuristic algorithm demonstrates that adaptive, uncertainty-driven sampling,combined with position-aware thresholding, can effectively balance mapping accuracy and resource efficiency. It produces an accurate environmental representation while minimizing redundant measurements,

laying a strong foundation for future integration with more advanced policies such as deep reinforcement learning.

5.2. Coverage map creation - Real Measurements

An important outcome of the validation campaign is illustrated in the comparative coverage maps shown above. On the left, measurements obtained with the Samsung Galaxy S23 Ultra using G-NetTrack Lite are presented, while the right-hand figure shows the same route measured with the SorusBoxScan device. In both cases, the average Reference Signal Received Power (RSRP) is displayed using a color scale ranging from deep red for very weak signals (below –100 dBm) to dark green for strong coverage (better than –80 dBm). This visual comparison provides a clear indication of the consistency between the two measurement platforms under real mobility conditions.

The results reveal a strong correlation between the readings of the commercial smartphone and the SorusBoxScan device. Segments of the route with weak coverage, particularly around the town of Parla and the southern sections near the tunnel areas, appear as red or orange in both maps, confirming the reliability of the SorusBoxScan platform in detecting low-signal regions. Likewise, areas with consistently good coverage, where RSRP remains above –85 dBm, are captured in green by both devices. The visual alignment across the entire measurement route demonstrates that the SorusBoxScan system can replicate the behavior of a professional-grade tool, even when deployed in a lightweight and flexible setup.

A closer inspection of the maps also highlights the benefits of using a dedicated measurement platform. The SorusBoxScan visualization on the right shows slightly smoother transitions between coverage categories, with less sporadic fluctuations compared to the smartphone data. This suggests that the SorusBoxScan's measurement pipeline, combined with its data buffering and publishing architecture, is capable of producing stable results even in areas with rapid coverage changes. Such stability is particularly important for long-duration campaigns or when measurements are used to build aggregate coverage models.

Overall, this comparison provides strong evidence of the validity and practicality of the SorusBoxScan system. By producing results that match those of a commercial smartphone application while offering additional flexibility in data handling and integration, the platform proves its value as a research and monitoring tool. The ability to capture and visualize coverage along real-world routes ensures that the device can be confidently used for advanced studies in mobile network performance, coverage optimization, and UAV-based measurement campaigns.

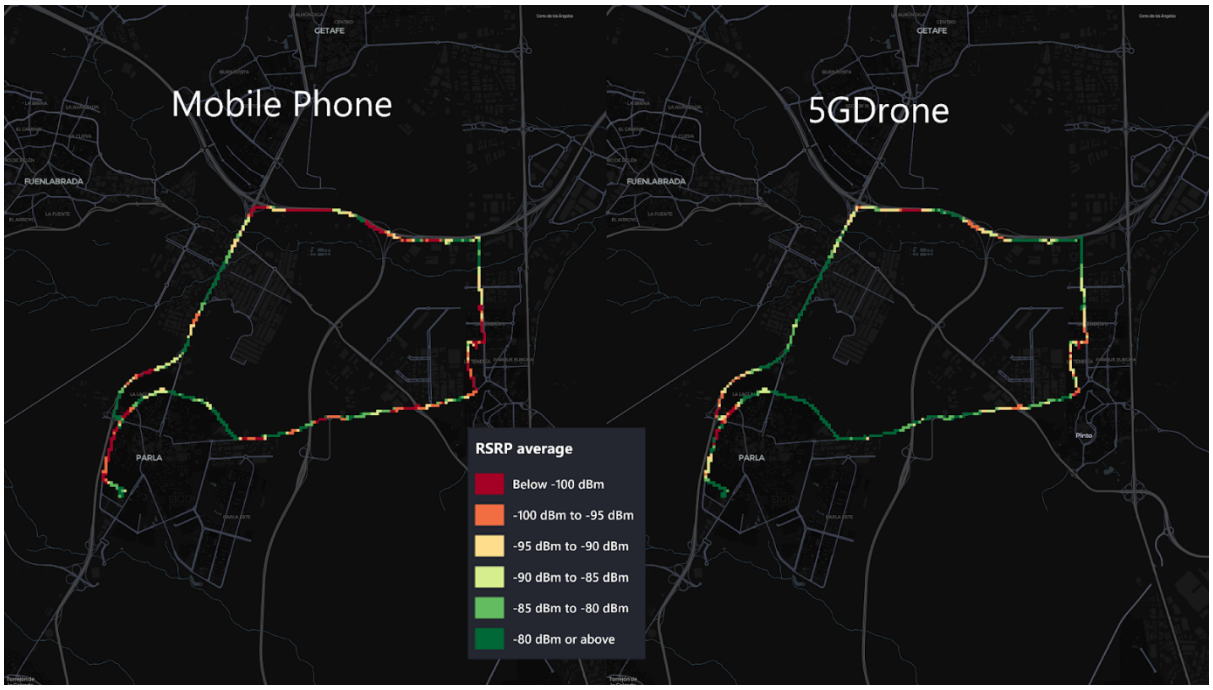

FIGURE 10: COMPARISON OF COVERAGE MAPS GENERATED BY A COMMERCIAL SMARTPHONE (LEFT) AND THE SORUSBOXSCAN DEVICE (RIGHT), SHOWING AVERAGE RSRP VALUES ALONG THE SAME ROUTE

6. Summary and Conclusions

This deliverable presented a UAV-based framework for efficiently creating wireless network coverage maps, focusing on its role in B5G/6G network management. The system integrates UAV-based data collection, probabilistic modeling using Gaussian Processes (GPs), and adaptive control strategies, forming a closed-loop workflow that iteratively improves coverage predictions while minimizing the number of measurements required.

We began by outlining the motivation for intelligent coverage mapping in emerging wireless networks, where traditional static mapping approaches are inadequate due to increasing network complexity, diverse service requirements, and dynamic propagation environments. UAVs offer a flexible and cost-effective platform for collecting measurements across challenging or wide-area terrains, while probabilistic models such as GPs provide a principled way to interpolate sparse measurements and quantify uncertainty in the predicted signal maps.

The challenges inherent to this task—such as collision avoidance, latency constraints, and efficient path planning—were discussed, emphasizing the need for autonomous strategies capable of balancing exploration and energy constraints. These challenges informed the formal definition of the problem, which was modeled as an iterative exploration task over a discrete grid (or graph) where each UAV action involves selecting a measurement point to reduce uncertainty in the coverage map.

To address this, we implemented two collection process models:

- Heuristic algorithm: A position-aware, uncertainty-driven approach that adjusts sampling thresholds based on node connectivity (corner, edge, inner) and considers both immediate and second-neighbor uncertainties. This method efficiently reduces uncertainty with fewer measurements compared to exhaustive traversal.
- Reinforcement Learning (RL): Although not evaluated within this deliverable due to its high computational cost and training requirements, RL was proposed as a future extension capable of learning environment-specific strategies that optimize exploration-exploitation trade-offs beyond heuristic rules.

The Gaussian Process-based modeling component was detailed for both one- and two-dimensional cases, demonstrating how GPs interpolate signal strength between sampled locations while maintaining a rigorous estimate of uncertainty. This capability is critical for guiding UAV exploration, as it identifies where additional measurements are most valuable.

In the evaluation, we demonstrated the system's effectiveness in progressively refining coverage maps through UAV-based sampling. The heuristic algorithm successfully prioritized high-uncertainty regions, achieving accurate reconstruction of the underlying environment with significantly fewer measurements than exhaustive approaches. The results confirmed the viability of combining UAV mobility with GPs for efficient network coverage mapping and established a strong baseline for future RL-driven strategies.

Finally, while this deliverable focused on simulation-based validation, the framework is designed to extend naturally to real-world wireless deployments. Its generalization to graph-based environments also opens opportunities for coverage mapping in structured urban areas, rugged terrains, and indoor industrial settings, where grid discretization may be impractical.

In conclusion, this work demonstrates that UAVs, coupled with probabilistic modeling and adaptive sampling strategies, can deliver intelligent, resource-efficient network coverage mapping—an essential capability for B5G/6G networks. Future work will involve integrating reinforcement learning for dynamic policy optimization, incorporating collision avoidance and latency-aware control mechanisms, and validating the framework with real-world UAV flights and measurement data. Together, these steps will move this approach toward deployment-ready solutions that support predictive, self-optimizing network management in next-generation wireless systems.

7. References

[1] Burke, P.J. Demonstration and application of diffusive and ballistic wave propagation for drone-to-ground and drone-to-drone wireless communications. Sci Rep 10, 14782 (2020). https://doi.org/10.1038/s41598-020-71733-0

[2] Ivanov A, Muhammad B, Tonchev K, Mihovska A, Poulkov V. UAV-Based Volumetric Measurements toward Radio Environment Map Construction and Analysis. Sensors (Basel). 2022 Dec 11;22(24):9705. doi: 10.3390/s22249705. PMID: 36560074; PMCID: PMC9781936.

[3] Warczek, J.; Kozuba, J.; Marcisz, M.; Pamuła, W.; Dyl, K. Research on Mobile Network Parameters Using Unmanned Aerial Vehicles. Sensors 2024, 24, 5526. https://doi.org/10.3390/s24175526

[4] V. Platzgummer, V. Raida, G. Krainz, P. Svoboda, M. Lerch and M. Rupp, "UAV-Based Coverage Measurement Method for 5G," 2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall), Honolulu, HI, USA, 2019, pp. 1-6, doi: 10.1109/VTCFall.2019.8891252.

[5] Y. Zeng, X. Xu, S. Jin and R. Zhang, "Simultaneous Navigation and Radio Mapping for Cellular-Connected UAV With Deep Reinforcement Learning," in IEEE Transactions on Wireless Communications, vol. 20, no. 7, pp. 4205-4220, July 2021, doi: 10.1109/TWC.2021.3056573

[6] D. Lee, O. Ozdemir, R. Asokan and I. Guvenc, "Analysis and Prediction of Coverage and Channel Rank for UAV Networks in Rural Scenarios With Foliage," in IEEE Open Journal of Vehicular Technology, vol. 6, pp. 1943-1962, 2025, doi: 10.1109/OJVT.2025.3585537.

[7] Liu, X.; Zhong, W.; Wang, X.; Duan, H.; Fan, Z.; Jin, H.; Huang, Y.; Lin, Z. Deep Reinforcement Learning-Based 3D Trajectory Planning for Cellular Connected UAV. Drones 2024, 8, 199. https://doi.org/10.3390/drones8050199

[8] M. Theile, H. Bayerlein, R. Nai, D. Gesbert and M. Caccamo, "UAV Coverage Path Planning under Varying Power Constraints using Deep Reinforcement Learning," 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, NV, USA, 2020, pp. 1444-1449, doi: 10.1109/IROS45743.2020.9340934.

[9] Rahman, M., Guvenc, I., Ramirez, D., & Wong, C. W. (2025). UAV-Assisted Coverage Hole Detection Using Reinforcement Learning in Urban Cellular Networks. arXiv preprint arXiv:2503.06494.

[10] B. Galkin, B. Omoniwa and I. Dusparic, "Multi-Agent Deep Reinforcement Learning For Optimising Energy Efficiency of Fixed-Wing UAV Cellular Access Points," ICC 2022 - IEEE International Conference on Communications, Seoul, Korea, Republic of, 2022, pp. 1-6, doi: 10.1109/ICC45855.2022.9838809.

[11] B. Omoniwa, B. Galkin, and I. Dusparic, "Communication-enabled deep reinforcement learning to optimise energy-efficiency in UAV-assisted networks," Vehicular Communications, vol. 43, p. 100640, 2023, doi: 10.1016/j.vehcom.2023.100640.

[12] Mundlamuri, R., Esrafilian, O., Gangula, R., Kharade, R., Roux, C., Kaltenberger, F., ... & Gesbert, D. (2023). Integrated access and backhaul in 5g with aerial distributed unit using openairinterface. arXiv preprint arXiv:2305.05983.

[13] Moore, J., Abdalla, A. S., Ueltschey, C., & Marojevic, V. (2024). Prototyping o-ran enabled uav experimentation for the aerpaw testbed. arXiv preprint arXiv:2411.04027.

[14] Y. -C. Zhi, Y. C. Ng and X. Dong, "Gaussian Processes on Graphs Via Spectral Kernel Learning," in IEEE Transactions on Signal and Information Processing over Networks, vol. 9, pp. 304-314, 2023, doi: 10.1109/TSIPN.2023.3265160.

[15] M. Stadler, B. Charpentier, S. Geisler, D. Zügner, and S. Günnemann, "Graph Posterior Network: Bayesian Predictive Uncertainty for Node Classification," in Advances in Neural Information Processing Systems (NeurIPS), vol. 34, pp. 18033–18048, 2021

[16] Alfaia, R.D.; Souto, A.V.d.F.; Cardoso, E.H.S.; Araújo, J.P.L.d.; Francês, C.R.L. Resource Management in 5G Networks Assisted by UAV Base Stations: Machine Learning for Overloade

Annex – installation of the SorusBoxScan

This guide explains how to set up a SorusBoxScan device so that it can connect to a mobile network and stream geo-tagged radio metrics to an MQTT broker. While the original documentation provides concise installation steps, here we expand on those instructions by explaining why each step matters, adding troubleshooting advice, and describing best practices for field deployments. Particular attention is given to the Waveshare SIM8200EA-M2 5G HAT, which is powered by the SIMCom SIM8200EA-M2 5G NR module. Understanding the characteristics of this hardware is essential because the quality of connectivity and measurement reliability depend heavily on how the device is configured, powered, and mounted.

Prerequisites

Before beginning the installation process, ensure that the hardware is properly assembled and all connections are secure. The SIM card must be inserted into the HAT, but keep in mind that SIMs with PIN codes enabled are not supported by the software stack; you should disable the PIN beforehand. The operating system should be Ubuntu, installed either on the drone's onboard computer or on a companion single-board system, and you should have root access either directly or via SSH. Finally, it is critical to connect the supplied cellular and GNSS antennas and position them so that they are not obstructed by metal surfaces or nearby electronics, as this can cause severe RF interference and loss of signal quality.

Connecting to the Mobile Network

In order to bring up the mobile data connection, the first step is to install the required QMI userspace tools. On Ubuntu, this is done with a simple `apt update` followed by the installation of `libqmi-utils` and `udhcpc`. Once these are in place, you should create a small helper script, often named `startnet.sh`, which initializes the modem, brings up the WWAN interface, and requests an IP address. The script typically calls `qmi-network /dev/cdc-wdm0 start`, followed by commands to activate the interface (usually `wwan0`) and acquire an IP lease.

Some operators require a specific APN to be set in `/etc/qmi-network.conf`. For instance, Movistar Spain requires the APN string `telefonica.es`. If your operator has different requirements, update the configuration file accordingly. It is also worth noting that `/dev/cdc-wdm0` is the QMI control device used to communicate with the modem, while `wwan0` is the associated network interface. Depending on how your system enumerates the hardware, the interface may appear as `wwan1` or another name. For persistent connectivity, especially when running unattended on a drone, you may want to create a systemd unit to launch this script automatically after the USB device is detected.

Installing the Control Software

The control software is distributed as a Snap package. This Python-based daemon communicates with the HAT through AT commands and converts the raw data into structured, geo-tagged metrics

for MQTT. Two commands are central: `AT+CGPSINFO=1`, which activates continuous GNSS streaming for location and movement data, and `AT+CPSI=1`, which reports the current radio access technology, frequency band, and signal quality indicators.

To install the software, download the `.snap` release from GitHub and install it with the Snap command line, using the `--dangerous` flag because the package is unsigned. If you want to make custom modifications, you can clone the repository, install Snapcraft and Git, and build your own `.snap` file. This is useful if you wish to add new AT commands, change reporting intervals, or adapt the software to your broker's data model.

Configuring MQTT Broker Settings

Once the snap is installed, you must configure the service so that it points to your MQTT broker. This is done with the `snap set` command, where you specify the broker's host, port, username, and password. You can enter each parameter separately or combine them in a single command. After updating the configuration, restart the daemon so that the changes take effect. From this point forward, all collected metrics will be published to the specified broker under the configured credentials.

If the device does not connect as expected, several steps can help diagnose the problem. Increasing the logging level to `DEBUG` and restarting the service will generate detailed logs, which you can follow in real time using `journalctl`. If no IP address is assigned, the most common causes are a missing or incorrect APN, a SIM card not being recognized, or the QMI device node not appearing in `/dev`. In such cases, check the output of `lsusb` to verify that the SIMCom module is enumerated properly. For GPS issues, reposition the GNSS antenna with a clear view of the sky and double-check that it is connected to the correct port.

SIM8200EA-M2 5G HAT, Technical Notes

The SIM8200EA-M2 5G HAT is built around SIMCom's SIM8200EA-M2 module, which uses the Qualcomm Snapdragon X55 modem. The module follows the M.2 3052 form factor and exposes a USB 3.1 data interface to the host. When used with a Raspberry Pi 4 or 5, it is critical to connect the board to a native USB 3 port in order to achieve high throughput and stable operation.

The module supports a wide range of 5G NR, LTE, and WCDMA frequency bands, though availability depends on your regional variant and operator certification. It also integrates a powerful GNSS receiver capable of using multiple constellations simultaneously. Proper antenna placement is crucial: the GNSS antenna should be mounted separately from the cellular antennas by at least ten centimeters to avoid interference. The HAT typically breaks out multiple antenna ports to support MIMO and positioning, and only matched 50-ohm antennas should be used to prevent detuning or high VSWR.

Power requirements are another important consideration. Under 5G bursts, the module may draw over two amperes at five volts, which means that a high-quality power supply capable of at least

three amperes is strongly recommended. Insufficient power can cause USB resets or unexpected disconnections. For drones, this implies careful power budgeting, since both motors and radio hardware may peak at the same time. Thermal management is also necessary: under sustained data transfer, the module will generate significant heat, so airflow or heatsinking should be planned in advance.

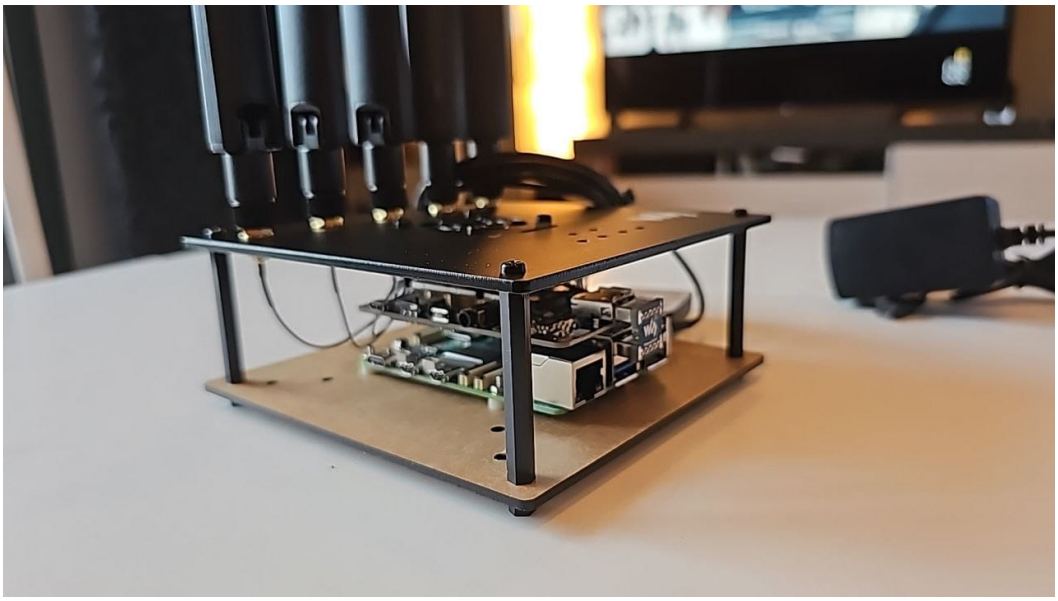

FIGURE 11: SORUSBOXSCAN MEASUREMENT DEVICE PROTOTYPE

From a deployment perspective, field practitioners should avoid cheap USB hubs and instead use direct, shielded USB 3 cables. The module's reset and flight mode pins allow recovery from unusual states, but in practice, a complete power cycle is sometimes the most effective method. When deploying on drones, antennas should be mounted vertically and kept away from carbon fiber frames or other conductive surfaces that can detune them. Short, low-loss coaxial cables with proper strain relief are preferred, especially in vibrating airborne environments. Finally, always confirm that your operating bands are legal in your region, and where possible, use antennas certified by the network operator.

Once operational, the HAT continuously streams GNSS data once per second and reports radio information through `+CPSI`. The software daemon merges these two data sources and generates MQTT messages where each radio measurement is tagged with a geographic position and timestamp. If the modem reports multiple carriers or bands simultaneously, the daemon emits multiple MQTT messages, each representing a distinct radio context but sharing the same spatial and temporal tags. This ensures that the dataset remains consistent and easy to correlate with coverage maps or flight telemetry.

Finally, when deploying a SorusBoxScan device, it is important to secure the data pipeline. Always configure MQTT with TLS, typically on port 8883, and assign unique credentials to each device rather than reusing a single account. Restrict publishing rights so that the device can only write to the necessary topics and rotate credentials periodically. Credentials should be stored only in the

Snap settings, never in plaintext scripts. To maintain security, update Ubuntu regularly and refresh installed Snaps so that both the system and application receive the latest patches.