ASCEND: Altitude Selection for High-quality Cellular Connectivity on Drones

Yongjae Yoo, Jihwan Suh, Dongrak Choi, Jeongyeup Paek, and Saewoong Bahk

***Abstract*—High-quality wireless communication is paramount to coordinating flights and missions effectively on a micro-UAV (a.k.a. drone). However, most cellular antennas are optimized for users on the ground and have not been planned for aerial devices. Many have studied cellular communication quality, but few works explore the impact of altitude during a flight. Through real-world experiments using an actual drone, we demonstrate significant connectivity dynamics in the air for both 4G LTE and 5G NR, and reveal that fixed-coordinate flights cannot maintain high quality connectivity in response to those dynamics. To address this problem, we present *ASCEND*, a reinforcement learning-based 3D altitude selection scheme that maintains high-quality connectivity during a flight over a planned 2D path without requiring prior training. We evaluate *ASCEND* at multiple real-world locations to demonstrate a notable increase in expected throughput and a reduction in the proportion of low-quality legs during a flight mission.**

***Index Terms*—Unmanned aerial vehicle (UAV), drone, cellular connectivity, real-world measurements, altitude selection**

I. Introduction

Advent of high-performance microprocessors has spurred significant advancements in micro-UAV (also known as drone) technology in recent years. Low-cost autopilot drones have become available that can operate by simply entering a 3D route without requiring real-time human intervention. Furthermore, advances in cellular communication in terms of both wide-area coverage and improved data rates have enabled videos to be streamed in real-time over cellular networks. These advancements have made new drone-based applications such as target tracking and mobile surveillance viable and promising [1]–[4]. They are receiving attention for the drones' mobility, deployment flexibility, wide-area coverage, and ability to track and follow [5]–[7].

Video surveillance using drones can be useful in a variety of situations, such as regular or occasional patrols of wide areas where installing fixed cameras would be costly and inefficient. Drones can also be deployed for rapid monitoring of new places, such as outdoor events, where the installation of surveillance equipment is difficult and time-restricted. Furthermore, the quality of cameras and cellular networks have become sufficient to stream high definition (HD) videos in real-time. Therefore, the idea of using drones for surveillance and patrol as in Fig. 1 became viable, and is no longer just a technology of the future.

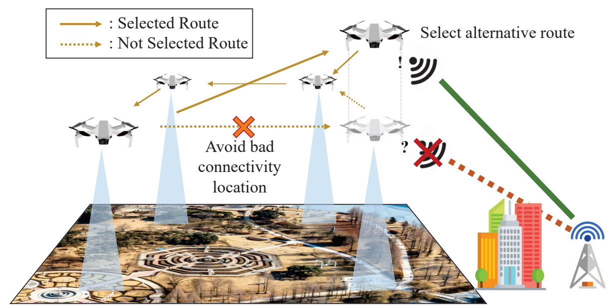

Fig. 1: Cellular-connected drone-based application example: surveillance and patrol.

However, the viability of such missions hinges critically on stable, high-throughput connectivity. For tasks that require high-resolution imagery or real-time monitoring, unpredictable signal drops triggered by altitude changes can corrupt data streams or lower video quality below usable thresholds, rendering the collected information unreliable for analysis or immediate action.

These challenges arises fundamentally because current cellular antennas are installed to serve users on the ground. They have not been planned nor investigated for use on aerial devices [17]. Furthermore, few works explore the impact of altitude on cellular connectivity of drones while moving, and even fewer use real measurements. Unreliable connectivity can lead to loss of control or interrupted video streams, particularly in unpredictable coverage areas. Through real experiments (§III), we demonstrate that the communication quality on a drone moving in the air has significant dynamics for both 4G LTE and 5G NR due to this reason[1]. We conducted extensive signal quality measurements using an actual drone at three distinct real-world locations. The results reveal that cellular connectivity in the air does not consistently improve or degrade with altitude, but rather exhibits irregular variations.

Based on our observations from real measurements, we find it critical for the drones to avoid communication shadow/gray zones during their flight, patrol, target tracking, and surveillance missions [1]–[4]. To address this, we propose *ASCEND*, a *reinforcement learning-based 3D altitude selection scheme*

Y. Yoo, J. Suh, D. Choi, and S. Bahk are with the Department of Electrical and Computer Engineering and INMC, Seoul National University, Seoul, Republic of Korea (email: {yjyoo,jhsuh,drchoi}@netlab.snu.ac.kr, sbahk@snu.ac.kr). J. Paek is with the Department of Computer Science and Engineering, Chung-Ang University, Seoul, Republic of Korea (email: jpaek@cau.ac.kr). S. Bahk and J. Paek are the co-corresponding authors.

This work was supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) under the ITRC(Information Technology Research Center) support program (IITP-2024-2021-0-02048), Next-generation Cloud-native Cellular Network Leadership Program(IITP-2025-RS-2024-00418784) funded by the Korea government (MSIT), and also by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2024-00359450).

[1]The 5G NR we refer to in this work is non-standalone (NSA). In the author's experimental environment, service providers offer NSA 5G NR only.

TABLE I: Comparison to existing research works based on key criteria

Research Work	Computational Complexity	Connectivity Consideration	3D Path Planning	Dataset Type	User-Defined Path
[8]	Low	X	X	Simulation only	X
[9]	High	X	O	Simulation only	X
[10]	Low	X	O	Simulation only	X
[11]	Low	X	O	Simulation only	X
[12]	High	O	O	Simulation only	X
[13]	High	O	O	Simulation only	X
[14]	High	O	O	Simulation only	X
[15]	Low	O	X	Simulation only	X
[16]	High	O	X	Real-world dataset	X
ASCEND	Low	O	O	Real-world dataset	O

(§IV) that maintains high-quality connectivity during a flight over a planned 2D path without requiring prior training. We discuss the characteristics of measurement data to justify the use of *reinforcement learning* (RL), and also analyze the impact of the reward design.

The contributions of this paper are as follows:

- We measure cellular signal quality in the air for both 4G LTE and 5G NR using a real drone, and analyze their characteristics as well as their relation to uplink throughput.
- We propose a novel 3D altitude-selection scheme, *ASCEND*, that allows a drone to autonomously optimize connectivity while on its flight mission. *ASCEND* is lightweight, and requires no prior learning.
- We implement *ASCEND* using a smartphone mounted on a drone as a proof-of-concept, and experiment at three distinct real-world locations to show a performance improvement of 17.87%∼23.1%.

The remainder of this paper is organized as follows. We review related work in §II, and analyze cellular connectivity in the air in §III. The design of *ASCEND* is presented in §IV, and §V evaluates *ASCEND*. Finally, §VI summarizes our work.

II. Related Work

While many studies have explored cellular connectivity [18]–[20], only a few have investigated connectivity with altitudes. We summarize the prior work as follows.

UAV's cellular connectivity has been studied under aerial base station and surveillance scenarios. Muzaffar *et al.* [21] observed how 5G connectivity changes as the location and altitude change on an UAV through real radio link measurements. Homayouni *et al.* [22] investigated cellular connectivity of an UAV at different altitudes and LTE bands, and analyzed how it affects performance of ground users. However, neither studies suggested a mechanism to autonomously adapt the altitude of the UAV during its flight. In addition, the studies measured cellular quality without a specific context for the application, and the trajectory of the UAV was limited to a straight-line-only, which is not suitable for most application purposes.

Lee *et al.* [23], Wen *et al.* [24], and Chakraborty *et al.* [25] have all studied cellular connectivity of drones in scenarios where a UAV is used as a BS. However, the first two studies conducted only simulations without actual drone flights. Chakraborty's research which implemented cellular networks using actual drones offers valuable insights for our study. Nevertheless, their emphasis on identifying UE locations and adapting coverage over large areas does not coincide with our study where the drone is transmitting high throughput data.

Video streaming via UAVs has been studied in a few prior works. Bertizzolo *et al.* [26] proposed a closed-loop control system for video streaming from UAVs on 5G Open-LAN architecture. The goal is to optimize UAV's location and its transmission direction by minimizing uplink interference. However, the approach is relatively abstract at a high-level without specific consideration for application requirements of video streaming. Naveed *et al.* [5] conducted video streaming experiments on LTE-connected UAVs to understand the relation between *reference signal received power (RSRP)*, throughput, and handovers under surveillance scenario. This work is most close to ours. We extend this work by considering the altitude change of the UAV, and design a RL based altitude-selection scheme that maintains high-quality connectivity under channel variation.

UAV path finding is another topic of interest related to our work. However, most studies in UAV path finding are inadequate in addressing the issue of UAV cellular connectivity in real-world environments. Some studies focus on drone path planning but fail to take into account connectivity while in flight [8]–[11]. Mardani *et al.* [15] investigated the problem of path planning considering cellular connectivity in UAV-based video streaming, and Li *et al.* [16] addressed communication constraints between UAVs and UGVs by proposing a memetic algorithm for path planning in cooperative systems. However, neither studies considered altitude.

Another study by Wang *et al.* [12] investigates trajectory planning for UAV-assisted data collection using a Double Deep Q Network (DDQN) approach. Yin *et al.* [13] investigates intelligent trajectory design for UAV-aided communication using a reinforcement learning approach to maximize uplink sum rate. Lee *et al.* [14] explores deep reinforcement learning for UAV trajectory design, focusing on optimizing the UAV's path to serve mobile ground users. These three studies differ from our scenario in that the drone either collects data locally or functions as a base station, and they all share the common limitation of having high computational complexity.

Summary of comparisons to related works are presented in TABLE I. Most importantly, the majority of the previously mentioned studies rely solely on simulation data rather than real-world data. In the following section, we demonstrate the novelty of our work by showcasing the characteristics of a real-world dataset, which was collected from actual flights of drones and mobile phones. Moreover, prior studies allow

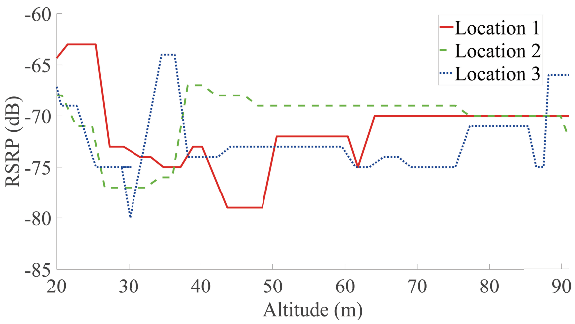

Fig. 2: RSRP according to altitude for three different locations, neither monotonic nor consistent.

drones to autonomously alter their flight paths, which imposes limitations on the scenarios. In most missions requiring cellular connectivity, the user needs to direct the drone to specific locations (e.g., surveillance using drones).

Our thorough literature review highlights a gap in solutions that allow user-defined missions while adaptively optimizing altitude for robust connectivity in real-time using lightweight, data-driven techniques. This motivates our proposal, ASCEND, designed specifically to address the challenge of maintaining high-quality cellular links for drones by dynamically selecting the optimal flight altitude during missions without prior training data.

III. Problem and Motivation

We first demonstrate the dynamics in cellular connectivity of drones to motivate our work. We then analyze the measurement data to make observations on how to design *ASCEND*.

A. Cellular Connectivity in the Air

To investigate cellular connectivity in the air, we mount a LTE/5G compatible smartphone on a drone and measure RSRP, RSSI, and SNR values[2], the commonly used cellular signal metrics that we can access on a smartphone [5], [21], [22], [27]–[29] at a rate of 2 samples per second (2 Hz). The three values have similar trends (more on this in §III-B), and thus here we show RSRP as an example representative metric for brevity. We describe the equipment we use and our detailed experimental setup later in the evaluation section (§V).

Preliminary Experiment 1: We examine cellular connectivity while varying the drone's altitude. The drone moves up and down from 20 meters to 90 meters above the ground at three different locations. During this vertical movement, the drone's coordinates in 2D (horizontal position) do not change. Fig. 2 plots the RSRP changes according to altitude. Contrary to intuitive expectations, the results were neither monotonic (increasing nor decreasing) nor constant. Furthermore, the altitude with the highest RSRP is not static. At location 1, the drone has the highest RSRP value at an altitude of around 25 meters. However at location 2, it is around 40 meters, and at location 3, peaks are at 35 and 90 meters.

[2]RSRP, RSSI, and SNR each stand for *Reference Signal Received Power*, *Received Signal Strength Indicator*, and *Signal to Noise Ratio*, respectively.

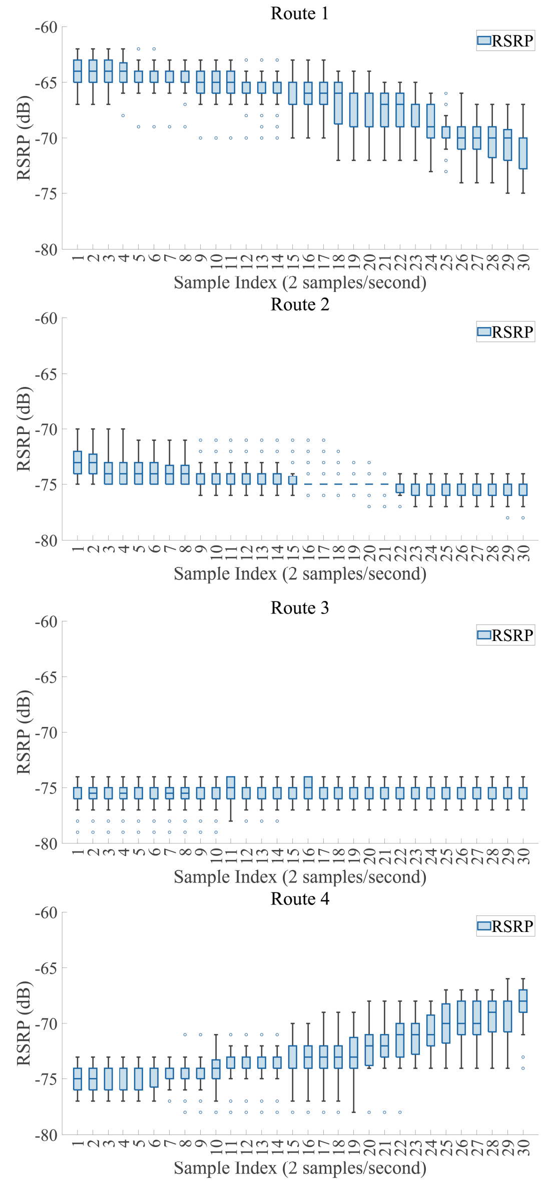

Fig. 3: Boxplot of RSRP on each of the four distinct routes, 50+ drone flights per route.

These observations motivate the need for a drone to change its altitude for better connectivity while performing its mission. Thus, our goal is to find an appropriate altitude to maintain good connectivity at run-time while a drone is on its flight mission. Then, the next question would be, "Can we find and use the best altitude for the mission based on historical data?". We answer this through our next experiment.

Preliminary Experiment 2: We configure a drone to fly static 3D routes repeatedly. We choose a total of 4 short routes, and the drone fly over 50 times for each route while measuring the cellular signal metrics. For example, Fig. 3 is the box plot of the RSRP obtained from this experiment (similar for RSSI and SNR). It shows that, although there are some trends in changes in RSRP in relation to location and altitude, RSRP varies significantly over time; i.e. over 50 iterations of the measurement shown by the long vertical boxes. Furthermore,

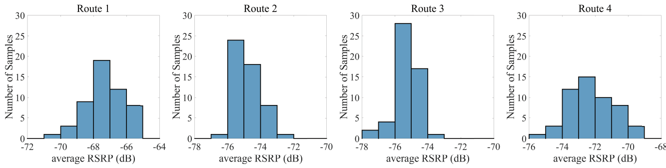

Fig. 4: Distribution of RSRP on each route

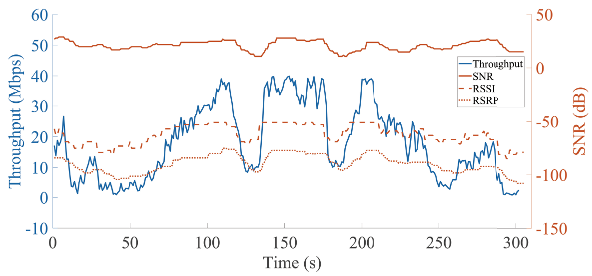

Fig. 5: Cellular network parameters (SNR, RSSI, RSRP) and their correlation with uplink throughput

the distributions of RSRP differ significantly per route as shown in Fig. 4. These results indicate that selecting the best route through an *a priori* measurement cannot guarantee a good path at later iterations. Therefore, our approach is to adopt machine learning to train a model that can predict the best path for the drone based on real-time data. Our intuition is that, in general, a dataset with such a probability distribution is well-trained by machine learning [30], [31].

B. Uplink Throughput and Parameter Selection

To continuously stream high-quality video from the drones, it is important to maintain sufficient uplink throughput. Ideally, if we know the uplink throughput at every location, we can easily and accurately select a suitable path for a mission. However, the achievable throughput at a specific location is a value that can only be known when large enough data is being (or has been) sent from that location. To design a mechanism that can select an appropriate altitude for the drone in the air during its mission, what would be need is a parameter that is readily available to infer or estimate the achievable throughput without having to try it out *a priori*. On the other hand, cellular signal metrics such as RSRP, RSSI, and SNR are already being measured periodically by the cellular system and available on end devices[3]. Therefore, we analyze these measurements to identify which parameter is appropriate, among those that can be obtained on recent smartphones for both LTE and 5G NR, for representing uplink throughput.

We install *iPerf3* on both a smartphone and a laboratory server, configuring it to transmit TCP uplink packets from the smartphone to the server. The reporting interval in our setup is 1 second, and the total test duration is approximately 300 seconds. We simultaneously measure the uplink throughput on the smartphone and use our custom mobile application to assess RSRP, RSSI, and SNR values. Fig. 5 plots the data where the X-axis is time, the left Y-axis is uplink throughput,

[3]Android PhoneStateListener API, https://developer.android.com/reference/android/telephony/PhoneStateListener

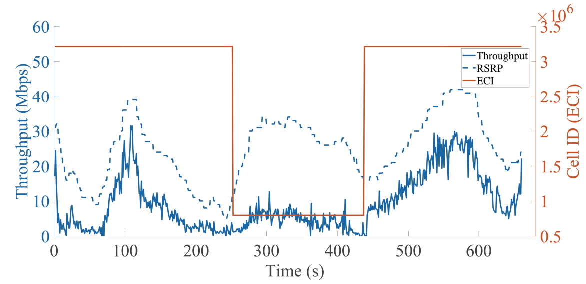

Fig. 6: RSRP and uplink throughput during handover. Handover occurs at 252 and 438 seconds, identified by cell-ID changes.

and the right Y-axis is for RSRP, RSSI, and SNR. The results show that RSSI, RSRP, and SNR all exhibit similar tendencies, moving up and down at similar times and locations. However, there is a slight difference in how well they mimic the uplink throughput. Among the three parameters, RSRP shows the highest correlation with uplink throughput, with a correlation score of 0.941.

Another consideration is handover. Since RSRP is the strength of a reference signal from a cellular base station, it may change abruptly when handover occurs. To verify this, we conduct an experiment on a path long enough for handover to occur (i.e. crosses the cell boundary), and measure RSRP and uplink throughput. Fig. 6 plots the result where the X-axis is time, left Y-axis is for scaled RSRP and throughput, and the right Y-axis is the current cellular network ID (ECI). The ECI values indicate that handover occurs at 252 and 438 seconds. Specifically, RSRP decreases as the drone approaches the cell edge, followed by a handover. Then, RSRP increases again as the drone moves into the new cell. Nevertheless, even with handover, RSRP follows uplink throughput well. Therefore, in this work, we decided to use RSRP as the main proxy for assessing cellular connectivity and estimating the achievable throughput to find an altitude for the drones to maintain stable and high-quality connectivity.

C. Preliminary Results

We summarize the results obtained from the experiments in Section §III as follows.

1) Among the common cellular signal metrics on a smartphone, RSRP correlates best with uplink throughput.
2) RSRP varies non-monotonically depending on altitude and time for same 2D coordinates, but
3) RSRP tends to follow a certain distribution at each location over long term.

For these reasons, our approach is to solve the problem using reinforcement learning, described in the next section. We argue that by cleverly adjusting the altitude, a UAV can achieve significantly better throughput throughout its flight.

IV. *ASCEND* Design

We present the design of *ASCEND* including its application requirements, components, reward, and their justifications motivated by the observations and insights obtained in §III.

This article has been accepted for publication in IEEE Transactions on Vehicular Technology. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TVT.2025.3594092

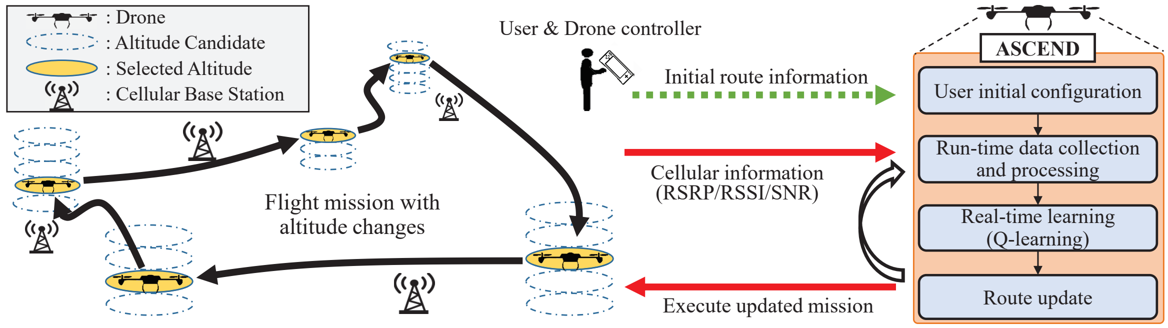

Fig. 7: System overview of *ASCEND*. *ASCEND* consists of four main components: *User initial configuration (waypoint)*, *run-time data collection and processing*, *real-time learning*, and *route update*. *ASCEND* selects an altitude for each waypoint among the altitude candidates while flying along the pre-configured 2D path.

A. Application Requirements

Consider a scenario where we need to quickly deploy a drone for continuous reconnaissance or patrol of a new site. The drone starts working at a new location immediately with no upfront data collection period. Accordingly, data for pre-training cannot be acquired, and the drone must start flying only with given 2D GPS coordinates. A solution must be real-time because the drone is constantly moving for its mission; the drone cannot stop-and-wait for next route to be calculated. Also, it should be light enough to be implemented on an embedded device such as the drone itself or a micro-processor/smartphone attached to the drone. Finally, 5G NR is not currently available everywhere in the world; there are many places where only LTE is used, or 5G and LTE are mixed with varying base station configurations. Accordingly, both cellular systems must be supported without modification[4]

Therefore, the goal is to design an autonomous altitude selection scheme for drones that can maintain high-quality cellular connectivity with no pre-training, has low delay, and is lightweight. From this point of view, commonly known behavioral cloning [32] and generative adversarial imitation learning techniques [33] require immoderate amount of pre-training data. In addition, the commonly used DQN [34], PPO [35], and SAC [36] methods require excessively large amount of resources (memory and computational cost) beyond the capabilities of mobile devices because they use a deep neural network. These techniques are not suitable for the scenarios we are targeting. To this end, *ASCEND* adopts Q-learning as a lightweight and low-latency reinforcement learning solution. We show the memory usage, computational cost, and the time it takes to accept new data and learn from a place in §V.

B. ASCEND Overview and its Components

Consider a scenario where a drone has a predefined route in 2D which consists of several waypoints. *ASCEND* completes the 3D path by appropriately determining the altitude of each waypoint in the drone's 2D route.

[4]Our approach is independent of LTE or 5G, and is compatible with both. We expect to have larger performance gains with stand-alone 5G because, compared to LTE, 5G signals have stronger directivity.

ASCEND consists of four main components as shown in Fig. 7: The user enters the initial setting values suitable for his/her task scenario (*User initial configuration*). Accordingly, the drone starts flying, while at the same time collecting and analyzing cellular information (*run-time data collection and processing*). *ASCEND* learns using the collected data (*real-time learning*), and determines the next altitude so that the system can operate with better connectivity (*route selection*).

User initial configuration: To start a mission, the user inputs a 2D flight route, and can also specify a range of altitude candidates (e.g. every 20 meters from 30 to 150 m) that can be selected by the drone. This initialization process needs to be done only once in advance, and can be altered later if necessary. *ASCEND* has high scalability because it is not affected by the number of waypoints or complexity of the path. In addition, it is not necessary to input any prior training data for learning.

Run-time data collection and processing: *ASCEND* uses a custom Android mobile application to obtain the RSRP, RSSI, and SNR values and the network type during its mission[5]. Then, according to the drone's flight path, *ASCEND* dissects the collected data between each waypoint (using altitude and time) and use it as data for subsequent learning.

Real-time learning: *ASCEND* learns the obtained data at run-time using Q-learning, and predicts the connectivity on the path. This process is lightweight, requires little memory, and has low latency in calculating the next route. The learning method of *ASCEND* using Q-table improves the connectivity of the drone by remembering the experienced route and searching for a new and better route.

Route update: As the result of the previous learning part, *ASCEND* calculates the next altitude in real-time. That is, the 3D path is completed by determining the altitude of each waypoint to pass next. Then, the drone flies automatically while changing its altitude for good connectivity.

C. Q-learning Design

Here we design the Q-learning for *ASCEND*.

[5]Android phones can read RSRP/RSSI/SNR values and their network type (LTE or 5G NR) as long as they satisfy Android API level 29 or higher.

Reward function of *ASCEND* is defined as Eq. (1).

$$Reward = avg_{RSRP} - std_{RSRP} + min_{RSRP}, \tag{1}$$

During its mission, the drone calculates the reward for each path segment between waypoints, using avg_{RSRP}, std_{RSRP}, and min_{RSRP}, which are the average, standard deviation, and minimum of the RSRP values encountered while traversing each segment.

- avg_{RSRP}: Intuitively, higher average RSRP is preferred since we can expect higher throughput (§III). Thus, a path with higher RSRP on average has a larger positive reward.
- std_{RSRP}: Larger deviation means that the drone will experience frequent throughput changes during its mission. This is unfavorable since the system may not respond correctly to the mission it was targeted for. Therefore, it takes a negative reward so that a path with lower deviation is preferred.
- min_{RSRP}: In addition, the system should avoid disconnections or too-low-quality connections at any instance of its mission even if the overall average quality is reasonable. Therefore, to maximize the min. RSRP, it is set as a positive reward.

Calibration between LTE and 5G data: As of 2023, 5G coverage is far below that of LTE, and thus a 5G user device often encounters places where it has to fall back to LTE. For this reason, *ASCEND* aims to operate well regardless of whether it is under LTE, 5G, or LTE & 5G mixed (often using a single base station in non-stand alone mode) coverage. In general, 5G NR has higher uplink throughput than LTE for the same RSRP value. Therefore, we calibrate the RSRP value in consideration of the network type to which the drone is connected while flying the corresponding route. The drone can make a decision on whether to opt for a route within the LTE network having a better RSRP or to follow a route within a 5G NR network, even if the RSRP is slightly inferior.

In order to correct the throughput gap between LTE and 5G with same RSRP value, we need to derive the relationship between RSRP and uplink throughput for LTE, 5G, and mixed networks. For this purpose, we measured 13,605 samples of throughput and RSRP for each network type on a Samsung Galaxy S22 smartphone. Our measurement data include samples with mobility at various altitudes and also in the presence of handover. Then, we approximate the measurement data as,

$$\begin{aligned} Thru_{LTE} &= \beta_0 + \beta_1 * RSRP_{LTE} + e \\ Thru_{5G} &= \beta_0{}' + \beta_1{}' * RSRP_{5G} + e' \end{aligned} \tag{2}$$

$Thru_{LTE}$ and $Thru_{5G}$ are the uplink throughput of LTE and 5G NR, respectively. Also, β_0, $\beta_0{}'$, β_1, and $\beta_1{}'$ are the coefficients of linear approximation. Lastly, e and e' are errors caused by the linear approximation. From our measurement data, the coefficients are calculated to be; $\beta_0 = 93.076$, $\beta_1 = 0.8873$, $\beta_0{}' = 145.64$, and $\beta_1{}' = 1.244$. Through this, we calibrate an RSRP value measured in LTE into an RSRP value in 5G which is expected to have the same throughput. *ASCEND* uses this calibrated value (converting LTE RSRP to 5G RSRP) for learning so that the same reward equation can be used regardless of the network type, i.e., LTE, 5G NR, or mixed.

Flexibility of *ASCEND*: *ASCEND* manages a Q-table of size, $\{state\}$ x $\{action\}$. In our scenario, we define the *state* as $\{waypoints\}$ x $\{altitude\_candidates\}$. This Q-table setting allows the user to easily set and change routes by simply adding or removing waypoints. Also, by modifying the altitude candidates, the user can limit the altitudes of the flight area where the drone is flying. Furthermore, depending on the scenario, it is even possible to learn using rewards other than the ones we have presented. For example, at the cost of allowing a little queuing delay, the user can put more weight on the average RSRP for higher throughput requirements such as transmitting a hologram or 360-degree camera data. Conversely, it is also possible to give more weight to the minimum RSRP when the user wishes to disallow receiving low-quality images such as when using the transmitted vision data for other machine learning purposes.

Example Scenario: We explain the operation of *ASCEND* using an example scenario. In this example, the user performs a surveillance mission within a square area. The user selects a total of four waypoints and sets altitude candidates of 50, 70, and 90 meters based on the terrain and altitude restrictions in their country. Therefore, we have a total of 12 states (4 waypoints $\times$ 3 altitudes). The action consists of three possible altitude choices, so the user's Q-table is configured as a 12×3 matrix.

Each entry in the Q-table represents the expected reward for taking a certain action (choosing the next waypoint's altitude) from the current state (the current waypoint and altitude of the drone). For instance, the value at position $(5, 1)$ of the Q-table indicates the expected reward if the drone is at the second waypoint at an altitude of 70 meters (5), and the next action is to choose an altitude of 50 meters (1).

When the drone actually chooses this action, it measures the RSRP during the movement to the next waypoint and calculates the throughput using the Eq. (2). The value is then plugged into the Eq. (1) to compute the reward, which is used to update the Q-table. As the mission progresses, the drone selects the best action from the Q-table or occasionally explores new actions, continuously updating the table to always search for the optimal path.

D. Discussion

Image quality due to altitude change. We use the surveillance scenario as an example to demonstrate the use of *ASCEND*, which changes the drone's altitude. A common concern is that the surveillance capability may deteriorate as the drone changes its altitude. To investigate this concern, we mount a camera on a drone and analyze the images taken at different altitudes (§V-D). Our analysis shows that there is no significant deterioration in image quality due to altitude change, within the legal ranges of altitudes that we use, with cameras comparable to today's latest smartphones. Detailed analysis results can be found in §V-D.

Energy/Battery consumption. Another concern is that the altitude change may result in additional battery consumption, impacting the energy efficiency of the drones. To investigate this question, we analyze the amount of battery consumed by the drone as it flies over various routes (§V-D). Our results show that, although the altitude change does have an effect on

(a) Argosdyne Aquilar drone **(b)** M16 drone controller **(c)** Battery and Velcro **(d)** Phone attached drone **(e)** Drone battery report from controller

Fig. 8: Experimental equipment

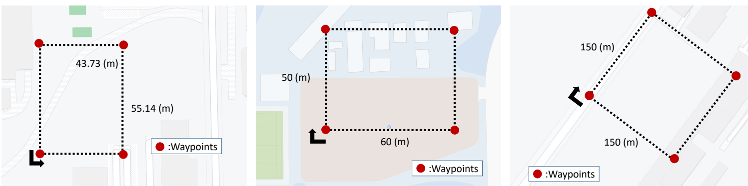

(a) Experiment location 1 **(b)** Experiment location 2 **(c)** Experiment location 3

Fig. 9: Experiment locations 1, 2, and 3 (Map details anonymized for review)

Fig. 10: Challenges from real-world experiments.

the battery consumption, the effect is relatively insignificant compared to the total path length of the 2D route and the environment in which the drone operates (e.g. wind). In other words, the total 2D path length and the wind factor dominates the energy consumption, and the altitude change is only a small fraction traded for improved uplink throughput. Detailed results can be found in §V-D.

V. EVALUATION

We evaluate *ASCEND* by comparing it against an *exhaustive search (ES)* and a *random jump (RJ)* schemes at three distinct real-world locations. We also discuss the implications of altitude change on image quality and battery consumption.

A. Experiment Setup

We use an *Argosdyne Aquilar* drone with an Android-based M16 controller (Figs. 8a and 8b), which can be configured to fly automatically along a pre-defined GPS route by specifying the coordinates of waypoints. During a flight, the drone moves only forward in a typical surveillance scenario. For example, to turn right, the drone changes its heading to the right and then proceeds forward. Consequently, the smartphone mounted on the drone rotates as well. The smartphone is always aligned with the front direction of the drone. Moreover, the drone flies at a constant speed of 5 m/s. These settings remain consistent across all experiments in this paper. We implement *ASCEND* and the cellular data collection application on a smartphone, and mount the phone on the drone (Fig. 8d). We use two smartphone models, the Galaxy note 9 and Galaxy S22, and develop the applications using Android API level 29.

During the drone's flight, a smartphone application attached to the drone records GPS, altitude, and signal strength information (such as RSRP). The results from the drone's repeated flight experiments are parsed using MATLAB, and the recorded altitude and rewards for each path are calculated to update the Q-table. In the Q-learning process, the learning rate is set to 0.1, and the probability of selecting a new path (epsilon) is linearly decayed from 1 by 0.01 at every 100th step.

We conduct experiments at three distinct locations: two urban and one suburban environments. The two urban environments are located close to each other and have many buildings that can affect the signals. In contrast, the suburban environment is more remote and does not have tall buildings or other reflective obstacles around.

First location. We select a rectangular flight path with four waypoints of size 43.73 x 55.14 meters within a university campus with some surrounding buildings[6]. We use four altitude candidates, 30, 50, 70, and 90 meters from the ground[7] for this experiment, but the user may choose any candidates within the legal range that suits their application requirements. We illustrate a schematic of the experimental site in Fig. 9a. At this location, measurement includes data for both LTE and 5G NR.

Second location. We evaluate *ASCEND* at multiple environments to show that the reward function we designed is not limited to a specific place and the scheme can be generalized. The second location is shown in Fig. 9b, and its size is 60 x 50 meters. Since the intention is to evaluate under a sufficiently different environment, we use a different smartphone (Galaxy Note 9), a different network type (LTE only), and a different TSP (Telecommunications Service Provider). In addition, we also change the altitude candidates to intervals of 10 meters at 30, 40, 50, and 60 meters from the ground. Note that *ASCEND* is *not* limited to four waypoints nor four altitude candidates; any number of waypoints and altitude candidates can be used; We use the same number for clear comparison, understanding, and presentation of the results.

Third location. This a suburban area far from the first two

[6]Choosing an experiment site require considerable effort. Our university campus is located in a large city where most areas are flight-prohibited zones. For safety reasons we had to avoid roads, and tree branches caught our drone multiple times as shown in Fig. 10.

[7]Going above 150 meters would require military permission at our site.

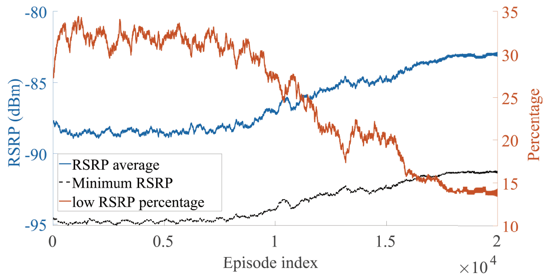

Fig. 11: Change of RSRP over episodes (Location 1)

locations, with a larger size of 150 x 150 meters as shown in Fig. 9c. At the third location, we use the same experimental settings as in the first location, except for the altitude candidates. We perform an extended-altitude scale experiment at 60, 90, and 120 meters from the ground at this location.

For comparison against *ASCEND*, we implement two baseline schemes: an *exhaustive search (ES)* scheme and a *random jump (RJ)* scheme; due to lack of comparative prior work with the same goal. In *ES*, the drone flies through all possible paths once in advance, then calculates and follows the *best path* among those paths. Thus, unlike *ASCEND*, *ES* needs to go through all routes once in advance. In addition, when using the *ES* scheme, the drone flies on a fixed route, making it impossible to change the route until all new possible routes have been examined again. We develop two different ways to choose a *best path* in the *ES* scheme. The first is $RSRP_{avg}$-*based ES* which picks the path with the highest RSRP average for better throughput. The second is *reward-based ES* which picks the path with the highest reward from Eq. (1). In the *RJ* scheme, the drone changes its altitude randomly when it transmits video with a quality lower than the reference quality predetermined by the user.

In total, we fly the drone more than ***1050*** times for the evaluation of *ASCEND*. The collected measurement data includes RSRP, altitude using a barometric pressure sensor, altitude using GPS, latitude, longitude, cellular base station ID (ECI), RSSI, RSRQ, CQI, and SNR, at a rate of 2 samples per second (2 Hz). The data and the software for the measurement have all been uploaded to GitHub, and will be made open-source for validation and public use.

B. *ASCEND Performance*

We evaluate the performance of *ASCEND* in several aspects. For convenience of explanation, drone test paths are named as follows: a vertex of a polygon becomes a waypoint, and the altitude of a waypoint is used as the name of that waypoint. For example, if the route name is {30, 50, 90, 30}, the drone flies along a 4-waypoint path, visiting waypoints at altitudes of 30, 50, 90, and 30 meters in order.

First, we check whether *ASCEND* learns correctly in the direction we intend. Fig. 11 plots the average RSRP, minimum RSRP, and the proportion of flight legs with low RSRP on the paths selected by *ASCEND* as learning progresses. The X-axis is the episode index, left Y-axis is for RSRP, and the right Y-axis shows the percentage of the interval on the path that experienced RSRP less than -95 dBm. As learning progresses, the average RSRP increases and the rate of the interval with low RSRP fell. Also, the lowest RSRP experienced in that episode has improved. As a result, the probability of a sudden cut-off in the corresponding path is significantly reduced. Also, the proportion of flight legs with low throughput is lowered. This shows that *ASCEND* learns the real-world data as intended in our reward design.

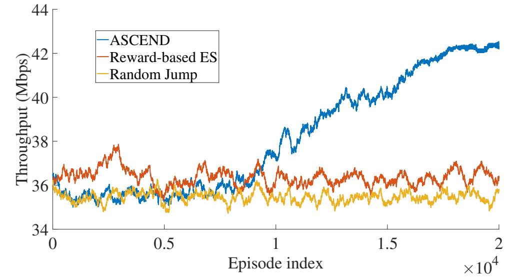

Fig. 12: Throughput comparison of *ASCEND*, *ES*, and *RJ*, trace-driven estimate based on measured RSRP (Location 1).

Next, we compare the performance of *ASCEND* with the baseline schemes *ES* and *RJ*. Fig. 12 plots the expected throughput of the three schemes based on the measured RSRP. It shows that the path selected by *ASCEND* continues to be superior to that of *RJ*, which means simply avoiding the low-performance routes reactively is insufficient and a better altitude should be predicted proactively. *ASCEND* does experience lower performance than *ES* in the early part of the episodes because *ES* has information about all routes in advance. However, *ASCEND* improves over time and has higher performance at later episodes because the *ES* uses fixed routes and does not respond well to changes in connectivity while *ASCEND* adapts using reinforcement learning. The average expected throughput of the first 5% section of *ASCEND* was 35.85 Mbps, but as learning progressed, the average expected throughput of the last 5% section was 42.26 Mbps, showing a performance improvement of about 17.87%. The result shows that if the drone continues to fly the same static 3D route as in *ES*, the temporal dynamics may worsen the performance and the opportunity to move to a better route is lost.

TABLE II presents the performance gains of *ASCEND* compared to other schemes. The first scheme is *ES* using $RSRP_{avg}$, which selects the route with the highest average RSRP after flying every route once. The second scheme is also *ES* but using our reward design (Eq. (1)), and the third is *RJ*. We calculate the estimated performance the drone experiences when flying each path using the measured path data and Eq. (2). We compare the average expected throughput and the ratio of expected video quality using the recommended bit rate table provided by YouTube [37].

First of all, *ASCEND* and *RJ* do not have a fixed flight path, whereas the $RSRP_{avg}$ based *ES* selects the path {50, 30, 30, 70}, and the *reward-based ES* selects the path {90, 30, 30, 70}. *ASCEND* achieved a higher throughput, 9.94% to 23.1% higher compared to the other schemes. It is worth noting that even though *ES* requires extensive pre-searching, *ASCEND* outperforms *ES* in terms of performance. This can be attributed to two main factors. Firstly, RSRP

		ASCEND	***$RSRP_{avg}$* based ES**	**Reward based ES**	**RJ**
Selected Route		No fixed route	50, 30, 30, 70	90, 30, 30, 70	No fixed route
Average throughput (Mbps)		42.36	38.53	36.96	34.41
Available video quality ratios (total 100%)	4K(2160p) @60 FPS (%)	9	0.5	0.25	0
	4K(2160p) @30 FPS (%)	37.75	35	28.25	31.84
	1440p @60 FPS (%)	53.25	64.5	71.5	68.16

TABLE II: Calculated throughput and video quality for each scheme

	Initial	1 month later		
Reward based ES: ***Route(ES_r)***	**Reward ranking**	**Reward ranking**	**Lowest RSRP**of ***Route(ES_r)***	**Lowest RSRP**of best route
90, 30, 30, 70	1^{st} / 256	193^{rd} / 256	-96 dBm	-81 dBm
$RSRP_{avg}$ based ES: ***Route(ES_{avg})***	***$RSRP_{avg}$* ranking**	***$RSRP_{avg}$* ranking**	***$RSRP_{avg}$* of Route(ES_{avg})**	***$RSRP_{avg}$*** of best route
50, 30, 30, 70	1^{st} / 256	59^{th} / 256	-84.31 dBm	-77.66 dBm

TABLE III: Performance of exhaustive search schemes over time

values collected during each drone flight using the *ES* method may not accurately represent the path due to RSRP fluctuations. Secondly, RSRP tendencies can change over time, which is discussed further in §V-C.

In the first experimental location, an urban scenario, the average throughput is relatively high, allowing for the transmission of videos at 1440p @60 FPS or higher regardless of the scheme used. However, the difference becomes evident in the distributable video quality ratios. Using other comparison schemes, the transmission of 4K (2160p) videos is possible only within a range of 28.5% to 35.5% during the mission flight, while with *ASCEND*, the transmission of 4K (2160p) is feasible for 47.75% of the time. This demonstrates that *ASCEND* not only increases the average throughput but also significantly impacts the ratio of video quality that end-users can experience. Moreover, if an user needs to transmit even higher quality videos, for example holograms or 360-degree camera footage in the near future, *ASCEND* will be well-suited to handle the higher throughput requirements. In summary, even when flying the same 2D surveillance path, *ASCEND* is capable of transmitting higher video quality by selecting the most appropriate altitudes.

Limitation of exhaustive search. When it comes to a long-term scenario, exhaustive search has another weak point. In the exhaustive search methods, the drone flies along a fixed route determined initially. However, if the monitoring is prolonged, the wireless channel may change due to changes in surrounding conditions. In this case, there is no way for the drone to switch to another route unless the drone performs a new exhaustive search. In order to confirm this, we experimented again at the same location approximately *a month later* and re-examined all possible flight paths.

TABLE III shows the performance of exhaustive search schemes over one month stretch. The route $\{90, 30, 30, 70\}$ previously selected by *reward-based ES* (*Route(ES_r)*) is no longer the best, and in fact exhibits the 193^{rd} reward out of a total of 256 possible routes. This means the drone is using one of the worst paths available. Furthermore, the lowest RSRP (the larger the better) experienced during *Route(ES_r)* is -96 dBm while there exists a route that has a lowest RSRP value of -81 dBm among the routes in our 1-month-later measurement. Meanwhile, the route $\{50, 30, 30, 70\}$ which had the highest average RSRP (*Route(ES_{avg})*), also could not avoid the performance drop. The drone flying along the *Route(ES_{avg})* achieves an average RSRP of -84.31 dBm, ranking only 59^{th} out of 256 possible routes, while the route with the best average RSRP among 256 routes achieves -77.66 dBm. Putting all together, we conclude that a flight with a fixed route through exhaustive search is greatly affected by environmental changes as the route cannot be changed unless a new exhaustive search is performed.

Computational cost. To profile the computation cost required to operate *ASCEND*, we run *ASCEND* using the measurement traces on an Intel i5-8500 3.00 GHz CPU without a GPU. As summarized in TABLE IV, the peak CPU usage was only 16.9% during training, required only 15.3 MB of RAM, and the average time to find the next path was 10.46 ms. We believe that *ASCEND* is lightweight and fast enough to run on the drones themselves.

C. Generalizability of ASCEND

Different settings. To show that *ASCEND* works well under different settings without modifying the reward design, we present Fig. 13 obtained from *location 2*. We differ not only the location, but also the phone, paths and their length, carrier and network type, as well as the altitude candidates to four at 30, 40, 50, and 60 meters. The results exhibit similar trends as with Fig. 11. As the episode progresses, the drone obtains a higher RSRP on average, and the minimum RSRP experienced by the drone during its mission improves. Consequently, the ratio of flights experiencing low RSRP is also reduced.

Different environment. We next move to a suburban environment at *location 3*. This is an open area with no large buildings, and there are no obstacles around that can cause meaningful reflections. Fig. 14 plots the expected throughput, which has similar tendency as the previous results from urban environments. Random jump or simply avoiding whenever a low RSRP is experienced is still insufficient to obtain high throughput. While *ASCEND* chooses a path that can achieve higher throughput as learning progresses, ES does not improve over time since it lacks the opportunity to choose a new path.

One notable characteristic in this environment is that while flying, the RSRP experiences a sudden and temporary drop even without any change in altitude. In particular, there are a number of samples with RSRP measured below -100 dBm, and

	ASCEND
Memory usage (MB)	15.3
Peak CPU usage (%)	16.9
GPU usage (GB)	Not needed
Inference time (ms)	10.46

TABLE IV: Resource usage footprint and inference time

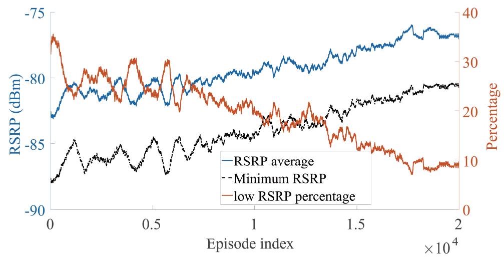

Fig. 13: *ASCEND* results at a different location with different settings (Location 2)

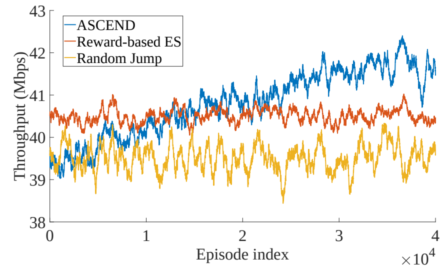

Fig. 14: Throughput comparison over time (Trace-driven estimate, Location 3)

Altitude	**30 m**	**50 m**	**70 m**	**90 m**
Snapshot	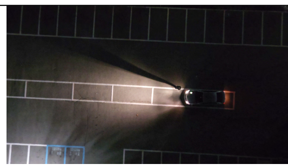	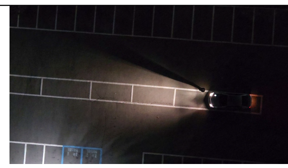	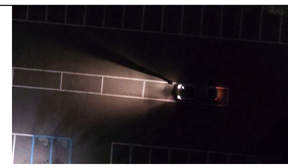	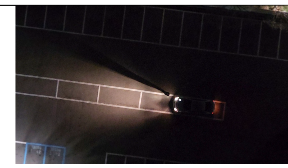
BRISQUE	46.58	48.42	51.54	49.47
NIQE	3.19	4.27	4.73	4.85
PIQE	38.92	39.02	47.26	50.24

TABLE V: Impact of altitude on image quality. Today's cameras with optical zoom are sufficient for surveillance purposes within the considered range of altitudes (i.e. 30∼90 m).

the lowest RSRP experienced during the mission is -139 dBm. In this case, the drone will experience a sharp drop in uplink throughput. Nevertheless, *ASCEND* makes the proportion of flight leg that experiences RSRP below -100 dBm to only 1.15%, and the average value of the lowest RSRP during the mission is -97.17 dBm when using *ASCEND*. This indicates that *ASCEND* is not over-fitted to a specific setting and can be used in other places without coefficient or reward changes.

We would like to emphasize that, the locations where we have conducted our experiments are highly populated metropolitan areas with one of the world's best LTE/5G coverage where it is rare to find outdoor areas with poor signal (within legal limits for UAVs). Furthermore, it is against regulations to fly UAVs in area with many buildings. We anticipate that our proposal will show significantly larger improvement in areas with sparser cell-tower placements and/or in area with many tall-buildings (where we are not allowed to fly currently).

Route Name	Action	Usage (mAh)
30, 30, 30, 30 (flat)	takeoff and landing	174
	horizontal flight	208
50, 50, 50, 50 (flat)	takeoff and landing	218
	horizontal flight	207
70, 70, 70, 70 (flat)	takeoff and landing	263
	horizontal flight	209
90, 90, 90, 90 (flat)	takeoff and landing	311
	horizontal flight	216
30, 50, 30, 30	whole mission	350
30, 70, 30, 30	whole mission	367
30, 90, 30, 30	whole mission	405
30, 70, 30, 70	whole mission	385
30, 90, 30, 90	whole mission	466

TABLE VI: Battery usage for different routes

Route Name	Usage (mAh)
30, 30, 30, 30	660
30, 30, 30, 30 (strong wind)	694
30, 70, 30, 70	675
30, 90, 30, 90	684

TABLE VII: Battery usage for larger routes

D. Impact of Altitude Change

Monitoring performance We analyze the camera's image quality change due to altitude and the resulting reconnaissance performance. For this purpose, we mount a Galaxy S22 smartphone on the same drone and record images at altitudes of 30, 50, 70, and 90 meters, respectively. In our target scenario, we assume that the drone is equipped with a typical smartphone-grade camera capable of optical zoom. TABLE V presents some of the images taken in the same area using only the different optical zoom levels. We experiment with drones at night to mimic a surveillance scenario.

We show how much the images are degraded through the following indicators: BRISQUE [38], NIQE [39], and PIQE. We randomly capture 10 images from the surveillance video and take average for each value using MATLAB. TABLE V confirms that each indicator does not change significantly depending on the altitude, showing that the photos taken by zooming in the same area have similar quality.

Battery consumption. We analyze two cases of battery usage in drones depending on the altitude. The first is the change in battery consumption that can occur when flying at an altitude different from the existing monitoring path; i.e. does flying at lower or higher altitudes change battery consumption? Next, we analyze the amount of additional battery consumption for the action of altitude changes. As shown in Fig. 8e, we measure the battery usage for each route using the cumulative battery usage provided by the drone controller.

First of all, TABLE VI shows that there is no significant difference in battery consumption regardless of the altitude when the drone is in horizontal-only flight. For the altitude changes, the difference in battery consumption between the route with the most frequent change and the route with the least change is only 116 mAh. On average, it shows a difference in battery consumption of about 23 mAh. That only accounts for a small percentage of the total mission and does not significantly reduce the drone's efficiency.

Finally the drone's battery consumption is ultimately determined by the *total flight time*, so the effect of altitude changes diminishes as the length of surveillance route increases. In order to show this, we expand the monitoring area into the size of 330 x 250 meters (4^{th} location). TABLE VII shows that the effect of altitude change is clearly reduced as the drone's surveillance route lengthens. In addition, we are able to confirm by experience that this level is even smaller than the battery consumption variance caused by *strong winds*. From these observations, we conclude that the extra battery consumption caused by the drone's altitude change is insignificant.

E. Discussion and Future Work

Our results demonstrate that *ASCEND* can achieve higher video quality in real-world applications. The duration of time during which only low-rate transmission can be supported is significantly reduced, leading to more consistent high-resolution streaming.

One limitation is that the size of the Q-table increases with each additional waypoint and altitude level, which extends the time required for convergence. As a trade-off for allowing the user to define their own path, this limitation becomes more pronounced as the complexity of the path increases. However, this is a common limitation that also applies to methods like exhaustive search, as the number of possible paths naturally increases. Nevertheless, the fundamental operational mechanism of *ASCEND* remains unaffected by this limitation.

Future research could explore methods like pruning techniques by incorporating additional information, such as base station locations, to reduce the computational burden and optimize performance. Additionally, this highlights a potential new area of research that could focus on maintaining user-defined paths while ensuring seamless communication.

VI. Conclusion

ASCEND is a technique for a drone to execute its mission on a given 2D path with good cellular connectivity by changing its altitude. We conducted extensive outdoor flight experiments with a real drone to identify the connectivity dynamics, and observed that cellular connectivity varies significantly with altitude, even for the same path. Evaluation results show that *ASCEND* can effectively adapt to dynamics in the air, resulting in a 17.87%$\sim$23.1% increase in expected throughput and a reduction in flight legs with low signal quality. *ASCEND* is light-weight, and gradually learns while performing its missions through reinforcement learning. We leave it as future work to evaluate the potential of this approach with degree-of-freedom other than altitude, and to assess its performance using real-time video streaming applications for surveillance.

References

[1] G. U. Sai Theja, M. S. Murari, M. F. Singha, R. Patgiri, and A. Choudhury, "A Survey on Surveillance Using Drones," in *Proceedings of the 14th International Conference on Contemporary Computing (IC3)*. ACM, 2022, p. 250–257.

[2] S. Kapania, D. Saini, S. Goyal, N. Thakur, R. Jain, and P. Nagrath, "Multi Object Tracking with UAVs Using Deep SORT and YOLOv3 RetinaNet Detection Framework," in *Proceedings of the 1st ACM Workshop on Autonomous and Intelligent Mobile Systems (AIMS)*, 2020.

[3] T. Hong, Q. Yang, P. Wang, J. Zhang, W. Sun, L. Tao, C. Fang, J. Cao, and B. Rong, "Multitarget Real-Time Tracking Algorithm for UAV IoT," *Wireless Communications & Mobile Computing*, vol. 2021, Jan 2021.

[4] H. Zhang, G. Wang, Z. Lei, and J.-N. Hwang, "Eye in the Sky: Drone-Based Object Tracking and 3D Localization," in *ACM International Conference on Multimedia (MM)*. ACM, 2019, p. 899–907.

[5] M. Naveed, S. Qazi, B. A. Khawaja, and M. Mustaqim, "Evaluation of video streaming capacity of UAVs with respect to channel variation in 4G-LTE Surveillance Architecture," in *International Conference on Information and Communication Technologies*, 2019, pp. 149–154.

[6] R. Shakeri, M. A. Al-Garadi, A. Badawy, A. Mohamed, T. Khattab, A. K. Al-Ali, K. A. Harras, and M. Guizani, "Design Challenges of Multi-UAV Systems in Cyber-Physical Applications: A Comprehensive Survey and Future Directions," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 4, pp. 3340–3385, 2019.

[7] L. Gupta, R. Jain, and G. Vaszkun, "Survey of Important Issues in UAV Communication Networks," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, pp. 1123–1152, 2016.

[8] X. Hou, F. Liu, R. Wang, and Y. Yu, "A UAV Dynamic Path Planning Algorithm," in *35th Youth Academic Annual Conference of Chinese Association of Automation (YAC)*, 2020, pp. 127–131.

[9] L. Ge, T. Shang, and C. Zhang, "A 3D Path Planning Algorithm for the UAV Based on the Improved Double-DqnAlgorithm," in *China Automation Congress (CAC)*, 2023, pp. 6412–6417.

[10] Y. Wu and K. H. Low, "An Adaptive Path Replanning Method for Coordinated Operations of Drone in Dynamic Urban Environments," *IEEE Systems Journal*, vol. 15, no. 3, pp. 4600–4611, 2021.

[11] V. Chiciudean and F. Oniga, "Pathfinding in a 3D Grid for UAV Navigation," in *IEEE 18th International Conference on Intelligent Computer Communication and Processing (ICCP)*, 2022, pp. 305–311.

[12] S. Wang, N. Qi, H. Jiang, M. Xiao, H. Liu, L. Jia, and D. Zhao, "Trajectory Planning for UAV-Assisted Data Collection in IoT Network: A Double Deep Q Network Approach," *Electronics*, vol. 13, no. 8, 2024.

[13] S. Yin, S. Zhao, Y. Zhao, and F. R. Yu, "Intelligent Trajectory Design in UAV-Aided Communications With Reinforcement Learning," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 8, pp. 8227–8231, 2019.

[14] W. Lee, Y. Jeon, T. Kim, and Y.-I. Kim, "Deep Reinforcement Learning for UAV Trajectory Design Considering Mobile Ground Users," *Sensors*, vol. 21, no. 24, 2021.

[15] A. Mardani, M. Chiaberge, and P. Giaccone, "Communication-Aware UAV Path Planning," *IEEE Access*, vol. 7, pp. 52 609–52 621, 2019.

[16] J. Li, T. Sun, X. Huang, L. Ma, Q. Lin, J. Chen, and V. C. M. Leung, "A Memetic Path Planning Algorithm for Unmanned Air/Ground Vehicle Cooperative Detection Systems," *IEEE Transactions on Automation Science and Engineering*, vol. 19, no. 4, pp. 2724–2737, 2022.

[17] S. Zhang, Y. Zeng, and R. Zhang, "Cellular-Enabled UAV Communication: A Connectivity-Constrained Trajectory Optimization Perspective," *IEEE Transactions on Communications*, vol. 67, no. 3, pp. 2580–2604, 2019.

[18] R. Ahmad, E. A. Sundararajan, N. E. Othman, and K. Arshad, "Handover in LTE-advanced wireless networks: state of art and survey of decision algorithm," *Telecommunication Systems*, vol. 66, no. 3, pp. 533–558, November 2017.

[19] H. Fourati, R. Maaloul, and L. Chaari, "A survey of 5G network systems: challenges and machine learning approaches," *International Journal of Machine Learning and Cybernetics*, vol. 12, pp. 385–431, Feb. 2021.

[20] T. Sylla, L. Mendiboure, S. Maaloul, H. Aniss, M. A. Chalouf, and S. Delbruel, "Multi-Connectivity for 5G Networks and Beyond: A Survey," *Sensors*, vol. 22, no. 19, 2022.

[21] R. Muzaffar, C. Raffelsberger, A. Fakhreddine, J. L. Luque, D. Emini, and C. Bettstetter, "First Experiments with a 5G-Connected Drone," in *ACM Workshop on Micro Aerial Vehicle Networks, Systems, and Applications (DroNet)*, 2020.

[22] S. Homayouni, M. Paier, C. Benischek, G. Pernjak, M. Leinwather, M. Reichelt, and C. Fuchsjäger, "On the Effect of Cellular-Connected Drones on Terrestrial Users: Field Trials," in *13th International Congress*

on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT), 2021, pp. 27–31.

[23] J. Lee and V. Friderikos, "Interference-aware path planning optimization for multiple UAVs in beyond 5G networks," *Journal of Communications and Networks*, vol. 24, no. 2, pp. 125–138, 2022.

[24] X. Wen, Y. Ruan, Y. Li, H. Xia, R. Zhang, C. Wang, W. Liu, and X. Jiang, "Improved genetic algorithm based 3-D deployment of UAVs," *Journal of Communications and Networks*, vol. 24, no. 2, pp. 223–231, 2022.

[25] A. Chakraborty, E. Chai, K. Sundaresan, A. Khojastepour, and S. Rangarajan, "SkyRAN: A Self-Organizing LTE RAN in the Sky," in *Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT '18)*. New York, NY, USA: Association for Computing Machinery, 2018, p. 280–292.

[26] L. Bertizzolo, T. X. Tran, J. Buczek, B. Balasubramanian, R. Jana, Y. Zhou, and T. Melodia, "Streaming From the Air : Enabling Drone-Sourced Video Streaming Applications on 5G Open-RAN Architectures," *IEEE Transactions on Mobile Computing*, vol. 22, no. 5, pp. 3004–3016, 2023.

[27] A. Fakhreddine, C. Bettstetter, S. Hayat, R. Muzaffar, and D. Emini, "Handover Challenges for Cellular-Connected Drones," in *Proceedings of the 5th Workshop on Micro Aerial Vehicle Networks, Systems, and Applications (DroNet'19)*, 2019, p. 9–14.

[28] M. Gharib, S. Nandadapu, and F. Afghah, "An Exhaustive Study of Using Commercial LTE Network for UAV Communication in Rural Areas," in *IEEE International Conference on Communications Workshops (ICC Workshops)*. Montreal, QC, Canada: IEEE, 2021, pp. 1–6.

[29] S. Homayouni, M. Paier, C. Benischek, G. Pernjak, M. Leinwather, M. Reichelt, and C. Fuchsjäger, "On the Feasibility of Cellular-Connected Drones in Existing 4G/5G Networks: Field Trials," in *IEEE 4th 5G World Forum (5GWF)*, 2021, pp. 287–292.

[30] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.

[31] A. Aksjonov and V. Kyrki, "A Safety-Critical Decision Making and Control Framework Combining Machine Learning and Rule-based Algorithms," 2022.

[32] F. Torabi, G. Warnell, and P. Stone, "Behavioral cloning from observation," arXiv preprint arXiv:1805.01954, 2018.

[33] J. Ho and S. Ermon, "Generative Adversarial Imitation Learning," *Advances in neural information processing systems*, vol. 29, 2016.

[34] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.

[35] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal Policy Optimization Algorithms," arXiv preprint arXiv:1707.06347, 2017.

[36] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor," in *International conference on machine learning*. PMLR, 2018, pp. 1861–1870.

[37] "YouTube bitrates and resolutions setting." [Online]. Available: https://support.google.com/youtube/answer/2853702

[38] A. Mittal, A. K. Moorthy, and A. C. Bovik, "No-Reference Image Quality Assessment in the Spatial Domain," *IEEE Transactions on Image Processing*, vol. 21, no. 12, pp. 4695–4708, 2012.

[39] A. Mittal, R. Soundararajan, and A. C. Bovik, "Making a "Completely Blind" Image Quality Analyzer," *IEEE Signal Processing Letters*, vol. 20, no. 3, pp. 209–212, 2013.

Yongjae Yoo is currently a Ph.D. candidate at the School of Electrical and Computer Engineering, Seoul National University, Seoul, Republic of Korea. He received a B.S. degree in Electrical and Computer Engineering from Seoul National University in 2019. His research focuses on enhancing wireless networks and IoT device through advanced software solutions, real-world cellular network measurements with drones, and applying machine learning to streamline these systems.

Jihwan Suh is currently a Ph.D. candidate at the School of Electrical and Computer Engineering, Seoul National University, Seoul, Republic of Korea. He received his B.E. degree in the field of Informatics and Mathematical Science from Kyoto University in 2013. His research interests are in the area of security, localization, and 5G with AI.

Dongrak Choi is currently Ph.D. candidate at the School of Electrical and Computer Engineering, Seoul National University, Seoul, Republic of Korea. He received a B.S. degree in Electrical and Electronics Engineering from Yonsei University in 2023. His research interests are in the area of UAV-assisted wireless networks, Internet of Things and 5G networks.

Jeongyeup Paek is a professor at the Department of Computer Science and Engineering, Chung-Ang University, Seoul, Republic of Korea. He received his B.S. degree from Seoul National University in 2003 and his M.S. degree from University of Southern California (USC) in 2005, both in Electrical Engineering. He then received his Ph.D. degree in Computer Science from USC in 2010. He worked at Deutsche Telekom Inc. R&D Labs USA as a research intern in 2010, and then joined Cisco Systems Inc. in 2011 where he was a Technical Leader in the Internet of Things Group. In 2014, he was with the Hongik University, Department of Computer Information Communication as an assistant professor.

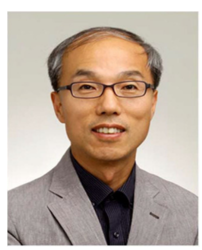

Saewoong Bahk received the B.S. and M.S. degrees in electrical engineering from Seoul National University (SNU), in 1984 and 1986, respectively, and the Ph.D. degree from the University of Pennsylvania, in 1991. He was with AT&T Bell Laboratories as a Member of Technical Staff, from 1991 to 1994, where he had worked on network management. From 2009 to 2011, he served as the Director of the Institute of New Media and Communications. He is currently a Professor at SNU. He has been leading many industrial projects on 3G/4G/5G and the IoT connectivity supported by Korean industry. He has published more than 300 technical articles and holds more than 100 patents. He is a member of the National Academy of Engineering of Korea (NAEK). He was a recipient of the KICS Haedong Scholar Award, in 2012. He was President of the Korean Institute of Communications and Information Sciences (KICS). He has been serving as Chief Information Officer (CIO) of SNU. He was General Chair of the IEEE WCNC 2020 , IEEE ICCE 2020 , and IEEE DySPAN 2018. He was Director of the Asia–Pacific Region of the IEEE ComSoc. He is an Editor of the IEEE Network Magazine and IEEE Transactions on Vehicular Technology. He was TPC Chair of the IEEE VTC-Spring 2014, and General Chair of JCCI 2015, Co-Editor-in-Chief of the Journal of Communications and Networks (JCN), and on the Editorial Board of Computer Networks Journal and the IEEE Tran. on Wireless Communications.