5G in the Sky: Uplink Throughput Measurement, Analysis and Enhancement

Yanbing Liu, *Student Member, IEEE*, Jingqi Huang, Ziyu Li, and Chunyi Peng, *Senior Member, IEEE*

***Abstract*—In this work, we present an in-depth study to measure, analyze and enhance aerial performance (here, uplink throughput) for drones flying in the low sky over two operational 5G networks in the US (AT&T and T-Mobile). Different from prior aerial 5G measurement studies, we have made three new endeavors. First, through extensive experiments in the low sky (below 120 m), we not only characterize aerial performance *observed* over operational 5G networks, but also quantitively assess performance potentials *not observed but missed* in the sky. We have several new findings that have not been reported before: higher 5G performance potentials are realized in the sky than on the ground (say, faster data speed in the sky); But surprisingly, more performance potentials are also missed in the sky (namely, 5G could have been even much faster but such potentials are not fully utilized in the sky). Second, we delve into root causes behind missed performance potentials and find that current 5G cell selection should take the blame despite the impacts of radio resource allocation in the underlying physical layer. Cell selection is designed for terrestrial scenarios and misses good 5G cells under aerial radio channel conditions. Third, we thus devise a data-driven solution called 5GAIR++ to patch cell selection in practice. 5GAIR++ is promising to pursue more 5G performance potentials in the low sky. We have validated its effectiveness over real-world traces with two applications of bulky file upload and video live streaming. Datasets and codes are released at [1].**

***Index Terms*—5G; Aerial 5G Performance; Uplink Throughput; Missed Performance Potential; 5GAIR++**

I. INTRODUCTION

Cellular-connected drones are gaining momentum with their emerging and thrilling uses such as aerial surveillance, construction inspection, environmental monitoring and conservation, post-disaster rescue, transport and logistics, and to name many [2]. Drones need to transfer a wide variety of data like videos, images, sensor data, commands, application-specific results and operation logs to their ground control systems and edge/cloud servers. Evidently, cellular networks offer an appealing communication option with tremendous advantages such as wide-area long-range coverage, seamless mobility support and quality data performance [2]–[4]. As illustrated in Fig. 1, cellular-connected drones can stay always connected wherever they fly in the low sky, say, below 400 ft (120 m) in the US allowed by FAA [5]. Aerial user equipment (UE), like terrestrial UE, performs handovers (HOs) to switch its serving cell from one to another (here, $C_X \rightarrow C_Z$, $C_Z \rightarrow C_W$), ensuring seamless connectivity on the fly (more background in §II).

Recent years have witnessed active efforts on measuring real-world performance of flying drones over cellular networks

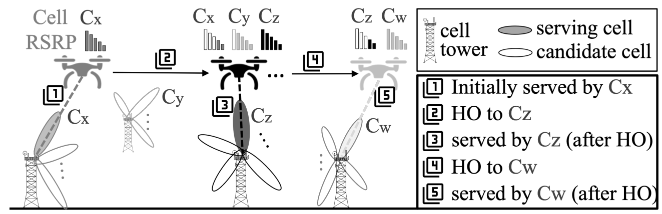

Fig. 1: Cellular-connected drones get "anytime, anywhere" connectivity in the low sky over the existing 5G/4G networks.

(e.g., [6]–[15]). These studies reported data throughput observed in their experiments but their results vary drastically and even contradict each other. For instance, the peak downlink rate of 600 – 700 Mbps was observed in [6], [7] as the test drone flied around one 5G cell tower, but [14], [15] saw no more than tens of Mbps over 5G. [7] found that HOs significantly impacted data performance but [14] believed that it was not the case; [14] reported that data speed slightly dropped at higher altitude but the opposite trend was observed in [15]. Similar issues also happened in prior measurements over 4G [8]–[10], [12], [13]. It is not hard to understand; All these studies run *piecemeal* measurements and thus observed performance depends on network deployment and runtime environments, which do vary a lot across distinct field trials.

However, it raises an important problem in generally understanding aerial performance over operational cellular networks, and gaining insights to enhance 5G support for drones in the low sky. More precisely, *How should we present and explain performance observed in real-world experiments so that the results are likely applicable to generic settings? How can we determine whether operational cellular networks perform well or not? If not, why do operational cellular networks fail to provide good performance as they can? How can we enhance data performance over the existing networks?*

In this work, we attempt to answer the above questions through our measurement study over operational 5G networks in the US. We focus on uplink data throughput because most drone applications (say, aerial surveillance) demand for high uplink data speed [16]. Different from previous measurement studies, we not only characterize aerial performance *observed* in our study but also investigate performance potentials which are *not observed but missed* in the wild. The core is to leverage extensive measurements to build performance reference so as to examine *relative* changes in data throughput instead of reporting *absolute* values only. By this means, we expect to reveal results and insights that are applicable to understand aerial 5G performance elsewhere. We characterize missed performance potentials by comparing performance observed

Y. Liu, J. Huang, Z. Li, C. Peng are with the Department of Computer Science, Purdue University, West Lafayette, IN 47907 USA. E-mail: {liu3098,huan1504,li4107,chunyi}@purdue.edu.

		Description	Figure(s)
5G Perf. (§III)	**F1**	Both operators (**A** and **T**) can offer much faster data speed in the sky than on the ground, but aerial 5G performance varies wildly with operator-specific variability at different altitudes.	Fig. 3, 4, 5
	F2	Two operators perform differently on the RAT level: **A** rarely uses 5G and **T** uses 5G a lot.	Fig. 7
	F3	For **T**, its aerial performance variance is primarily attributed to the use of distinct 5G cells.	Fig. 8, 9
	F4	Aerial performance is improved with more physical radio resources (more RBs and higher CRs).	Fig. 11
	F5	For **T**, performance variance is still dominated by the selection of serving cells because the range of resource block number is impacted by the serving cells.	Fig. 12
Missed(§IV)	**F6**	Good 5G cells are missed more often in the sky than on the ground.	Fig. 13, 14
	F7	The missing of good cells prevents more than half of potential throughput from being realized.	Fig. 15, 16
	F8	Operators miss good cells with distinct causes: no configuration (**A**,**T**) and poor HO decision(**T**).	Fig. 17
	F9	No 5G configuration is caused by **A**'s channel-specific policies and **T**'s preference on mid-band.	Fig. 19, 20a
	F10	For **T**, good cells may be missed due to radio-oriented HO decision, even if they are reported.	Fig. 20b, 21
Fix (§V)	**F11**	5GAIR++ fixes >50% of problematic HOs and doubles throughput in 30% – 50% of instances.	Fig. 24 – 27
	F12	It significantly improves experience of video live streaming (increases FPS and reduces stalls).	Fig. 28
	F13	The serving cells of aerial and terrestrial UE might be distinct (depending on freq bands).	Fig. 29
	F14	5GAIR++ yields even higher throughput gain in the presence of terrestrial UE in some cases.	Fig. 31

TABLE I: Summary of our main findings in this work (marked as F1 – F14).

at runtime (namely, performance potentials realized) with reference performance profiles. Interestingly, we see that performance variance at the same location is a good indicator to performance potentials realized and missed. Furthermore, we reason about performance potentials which are available but missed in reality. Inspired by our cause analysis, we design a patch solution to enhance aerial 5G performance over operational networks. Table I summarizes our main findings. We have made three main contributions in this work.

- **What does aerial 5G performance look like (§III)?** We have conducted a measurement study with two top-tier US operators (**A** and **T** for AT&T and T-Mobile afterwards) to quantitively characterized 5G performance in the low sky. We find that 5G is much faster in the sky than on the ground but its data throughput varies wildly for both operators. We conduct a top-down analysis on aerial performance variance at three levels: RAT, cell and physical layer. For **A**, it is mainly due to poor 5G usage at the RAT level but for **T**, performance variance is dominated by cell selection at the cell level, despite the impact of resource allocation at physical layer.

- **How much aerial performance potentials are missed and Why (§IV)?** We further characterize performance potentials missed in the sky. Good 5G cells are not often selected for drones in the sky, which prevents more than half of aerial performance potentials from being realized. We delve into an in-depth cause analysis and pinpoint critical factors causing missed good 5G cells, which varies with operators.

- **How can we improve aerial 5G performance (§V)?** Inspired by our findings, we propose 5GAIR++, a quick fix solution to enhance 5G performance for drones in the low sky. 5GAIR++ provides altitude-aware optimization at two key steps – configuration and decision making – through patching the legacy cell selection mechanism. Our evaluation over real world traces validates that 5GAIR++ effectively mitigates the missed performance issues, doubles uplink throughput in more than 30% – 50% of aerial instances and enhances smooth 4K video streaming in the air. Moreover, 5GAIR++ achieves even higher throughput gains in the presence of other ground users.

Release. Codes and datasets are are released at [1].

II. BACKGROUND

Radio access in cellular networks. In a cellular network, a cell is one *basic* unit to offer radio access to user equipment (UE). Each cell runs one radio access technology (RAT, say, 5G, 4G or 3G) over one contiguous spectrum frequency block (referred to as a frequency channel). It physically resides in a cell tower which accommodates a number of cells over distinct frequency channels and directional antenna (Fig. 1).

A serving cell is selected or re-selected via standard handover (HO) procedures, identical for terrestrial and aerial UEs. Each HO relies on radio quality (RSRP/RSRQ) measurements to decide if switching is needed, following five steps: *configuration, measurement, reporting, decision, and execution*. At the start, UE is served by one cell, which sends configured parameters to customize the subsequent measurement and reporting steps including cells to be measured and reporting criteria. These criteria are defined as reporting events (say, A1-A6, B1-B2) by comparing RSRP/RSRQ of the serving cell and candidate cells [17], [18]. UE then measures cells nearby and reports their measurements when the reporting criteria are met. Finally, the serving cell decides whether to execute a HO and switches to another if applicable.

5G features observed in this study. 5G follows the above procedure to establish and migrate radio access while adopting several advanced features observed in this study. First, 5G networks use both 5G and 4G (two RATs) to serve UE through dual connectivity [19]. US operators run 5G primarily in Non-Standalone (NSA) mode[1], where 4G acts as the master RAT and 5G offers secondary radio access [20]. Second, each RAT uses carrier aggregation to allow more than one serving cells [21]. As a result, UE is served by a set of serving cells, consisting of two cell groups (4G+5G) if 5G is used, otherwise one cell group over 4G (4G only). Each cell group consists of a primary cell (PCell) responsible for configuring and performing HOs, and several secondary cells (SCells).

[1]T-mobile used 5G NSA in this study when most experiments were done before May 2024 and now advance to 5G SA in the same test area. AT&T supports 5G NSA only till now.

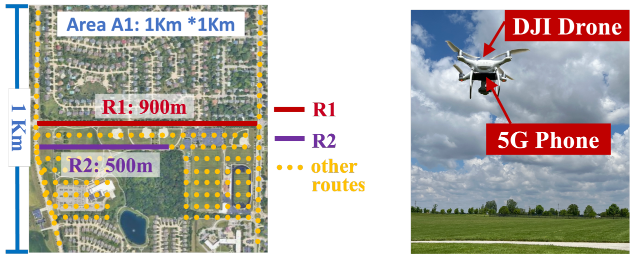

(a) Map (b) Aerial UE

OP.	Duration	Distance	#Cell (5G/all)	#HO	#Thput	#RSRP/Q
A	1,104 min	96 Km	26/1,881	702	0.6 M	2.8 M
T	3,167 min	343 Km	285/1,340	45,141	1.9 M	9.2 M

(c) Dataset statistics

Fig. 2: Experimental settings and dataset statistics.

III. What Does Aerial 5G Performance Look Like?

In this section, we present aerial 5G performance observed in our extensive experiments in the low sky, and then perform a breakdown analysis to understand performance variability.

A. Measurement Methodology and Datasets

We run extensive measurement experiments to characterize aerial 5G performance with two US operators (**A** and **T**) [2]. The experiments are primarily conducted in one 1 Km × 1 Km area in West Lafayette, IN (map in Fig. 2a). This test area is a typical residential zone fully covered by 5G with houses (top), parks and sport fields (center) and apartments (bottom). Since commodity drones with 5G connectivity are not available yet, we use a drone (DJI Phantom 4 Pro) carrying a 5G phone (Google Pixel 5) to measure aerial performance (Fig. 2b). We fly drones at different altitudes up to 120 m (allowed by FAA [5]). We conduct flying experiments along various flight routes including two selected routes (R1 and R2) and other routes on main roads and public areas. We also run driving experiments to measure performance on the ground. The moving speed in flying and driving experiments are 3-5 m/s and 10-15m/s, respectively. Such speed does not influence RSRP/RSRQ measurements which can be done on time before they are used for handover decisions. Moreover, we run driving experiment over the same route right after finishing flying experiments if applicable. By this way, we ensure similar experimental conditions for fair comparison while avoiding potential interference with multiple active phones. In this work, we focus on uplink data speed because many drone applications need to upload or stream heavy traffic (e.g., videos, images and sensor data). In all experiments, we test with bulky file upload (repeatedly uploading bulky files (50 MB) with a TCP connection to lab server) unless specified. Video live streaming is later used in the evaluation (§V-B).

We have conducted experiments sporadically from Nov 2023 to May 2024 with data logs over 4,181 minutes (Table 2c, **A**: 1,014 minutes and **T**: 3,167 minutes). In total, we have collected 45.8K HO instances with more than 2.5M throughput samples and 12M RSRP/RSRQ samples. We observe more than 3,200 cells including 311 cells over 5G. It indicates dense cell deployment with many candidate cells (tens of cells or more) which match with recent studies [22], [23].

[2] Verizon is not included in this study because we find that it had poor 5G coverage in the sky when most experiments were done (till May 2024).

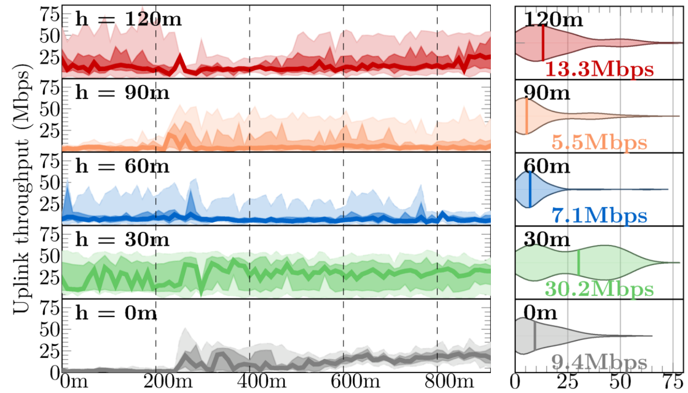

(a) The 5/25/50/75/95th-percentile of throughput (b) All

Fig. 3: Uplink throughput over route R1 at all test altitudes (A).

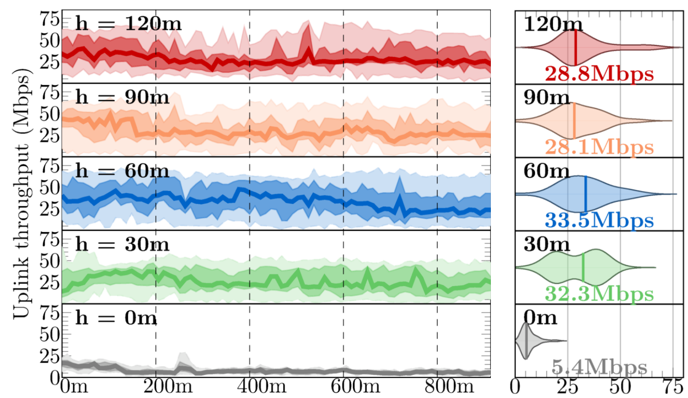

(a) The 5/25/50/75/95th-percentile of throughput (b) All

Fig. 4: Uplink throughput over route R1 at all test altitudes (T).

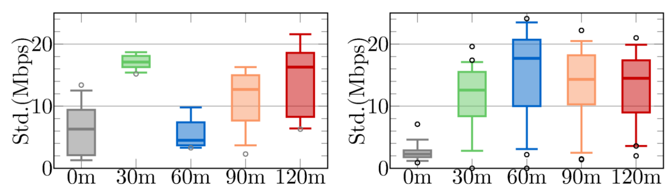

(a) Std. per altitude (**A**) (b) Std. per altitude (**T**)

Fig. 5: Boxplot of throughput standard derivation (std.) over R1.

B. A Glimpse of 5G Performance in the Low Sky

We first give a glimpse of operational 5G performance in the low sky using uplink data speed observed over one selected route R1. All main findings are consistently observed in the test area (§III-D, §IV-A). R1 is a 900 m route along one main road. We test with five different altitudes: 0m, 30m, 60m, 90m, 120m, each with extensive runs (>20) at different hours and/or different days. Fig. 3 and Fig. 4 plot uplink throughput observed with two operators (**A** and **T**). At each altitude, we show the 5/25/50/75/95th-percentile of uplink throughput per location (grid size: 10 m) and then give a violin plot using all throughput samples observed at all locations.

[F1] Both operators (A and T) can offer much faster data speed in the sky than on the ground, but aerial 5G performance varies wildly with operator-specific variability at different altitudes.

We have three observations. First, both operators can upload files much faster in the sky (here, at least at 30 m) than on the ground at the same location. Aerial UE achieves higher than 30 Mbps in the sky but the same phone gets <10 Mbps on the ground. Second, aerial 5G performance varies with operators at different altitudes. Over route R1, **T** provides good performance at all test altitudes in the sky but **A** offers high speed only at 30 m. The median throughput of **T** ranges from 28.1 Mbps to 32.4 Mbps at all altitudes in the sky, which is at least 5.2 times higher than on the ground (5.4 Mbps). In contrast, **A** cannot maintain high throughput at all altitudes. Its median throughput shrinks from 30.2 Mbps to several Mbps (5.5 – 13.3 Mbps) above 60 m, which is even worse or comparable to the throughput on the ground (9.4 Mbps). Third, we see that data throughput varies wildly in the sky. For both operators, aerial performance fluctuates from several Mbps (say, 5th-percentile) to several tens of Mbps (say, 95th-percentile) at most places. At each location grid, we calculate the standard deviation (std.) of uplink throughput and show the boxplot of throughput std. at all locations per altitude in Fig. 5. For both operators, the throughput std is larger than 12 Mbps at all altitudes except at 60m (**A**), which at least doubles the terrestrial one (2.3 – 6.3 Mbps).

We want to highlight distinct performance provided by two operators (**A** and **T**). In the sky, **T** is usually faster than **A** (see the median throughput at each location) but it does not mean that **A** cannot be that fast. We notice that **A** is able to offer high data speed in 5% of instances (its 95th-percentile throughput is much larger). It implies that **A** still has the potential to offer high throughput at most locations but such high throughput potentials are not often realized. We further compare the *relative* ratio of the throughput std. over the median throughput, and find that there exists huge improvement room for both operators. Especially, above 90 m, **A**'s relative ratios are much higher (90m: 13/5.5 = 236%, 120m: 16/13.3 = 120%) than **T** (90 m: 14.3/28.1 = 50.9%, 120m: 14.5/28.8 =50.3%). That is to say, for **A**, there exists a considerable performance gap between the potential throughput and the actual throughput realized at these altitudes.

There are two implications on 5G performance potentials. On one hand, higher performance potentials are realized in the sky. 5G can be fast and even faster than on the ground. On the other hand, such huge performance variance in the sky also implies considerable performance potentials might be missed in reality, as 5G networks offer low data speed at locations where faster speed is available. Different from our prior work on terrestrial performance potentials [24], [25], larger throughput variance in the sky implies that performance potentials are more likely missed in the sky than on the ground. We conduct a top-down analysis to understand why performance variance in the air is so high at three levels: RAT (§III-C), cell (§III-D), and physical radio resource (§III-E).

C. Performance Variance Between RATs

We first analyze the aerial performance variance at the RAT level. Surprisingly, we find that **A** and **T** perform differently.

[F2] Two operators perform differently on the RAT level: T uses 5G in the sky at most time but A rarely uses 5G, which is the primary source of its performance variance, missing high performance potentials offered by 5G.

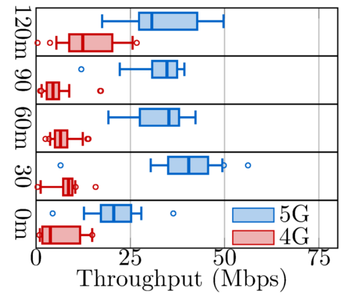

(a) Uplink throughput (**A**)

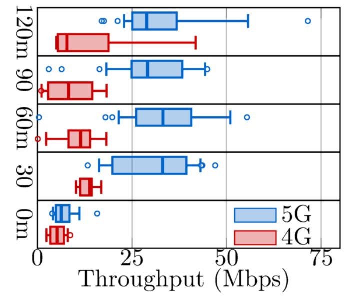

(b) Uplink throughput (**T**)

Fig. 6: RAT-level (5G/4G) data throughput over R1 (A and T).

	OP	Band	Ch. Freq	Ch. BW	#SCells	Usage
5G$_1$	**A**	n5 (low-band)	874 MHz	10 MHz	10	100%
5G$_2$		n41 (mid-band)	2600 MHz	100 MHz	55	76.2%
5G$_3$	**T**	n71 (low-band)	626 MHz	20 MHz	30	21.2%
5G$_4$		n71 (low-band)	649 MHz	20 MHz	1	2.6%

TABLE II: Main 5G channels observed in this study.

We first compare the performance at the RAT-level(say, 5G and 4G). There is no surprise that 5G is much faster than 4G in the sky for both operators, despite smaller bandwidth of 5G channels used by **A**. Table II lists four 5G channels used by serving cells (all as SCells) observed in this study. **A** uses one low-band channel (10 MHz) and **T** uses one mid-band channel (100 MHz) and two low-band 5G channels (both 20 MHz). Note that **A** does not deploy 5G over high-band channels (say, over mmWave bands, > 24 GHz) in this test area. This is different from our recent 5G measurement studies in big cities [22]. The use of distinct 5G channels impacts data performance, resulting in operator-specific issues for **A** and **T** as elaborated later.

Fig. 6a and Fig. 6b plot uplink throughput over 5G and 4G in **A** and **T**, respectively. Note that 5G is never used alone because both **A** and **T** run NSA 5G in our test area. All the performance is provided by a set of serving cells over 4G and 5G (4G+5G). We compare it with performance provided by 4G-only cellsets. Unsurprisingly, 5G (more precisely, 4G+5G) is much faster than 4G only for both operators. In all the aerial scenarios, 5G performs significantly better than 4G only. For **A**, 5G boosts the median throughput by 2.5×–7.7× from 5–12 Mbps to 31–41 Mbps. For **T**, data throughput rises from 8–14 Mbps to 29–33 Mbps with 5G, which is even slightly lower than **A**. It is anti-intuitive at the first glimpse because the bandwidth of one 5G channel used by **A** (10 MHz) is much smaller than those by **T** (20–100 MHz, see Table II). Despite smaller channel bandwidth, **A** is able to offer comparable or even slightly higher data speed than **T** in the low sky. We later explain that it is attributed to distinct physical-layer radio resource allocation between **A** and **T** (§III-E). These results imply that 5G cells are the primary contributor to high speed in the air, and we thus focus on the performance impacts of 5G cells afterwards unless specified. Moreover, we notice an exception on the ground, where 5G does not improve performance too much. For both operators, their 5G performance on the ground is significantly lower than the one in the air. For **T**, 5G even performs much worse and even

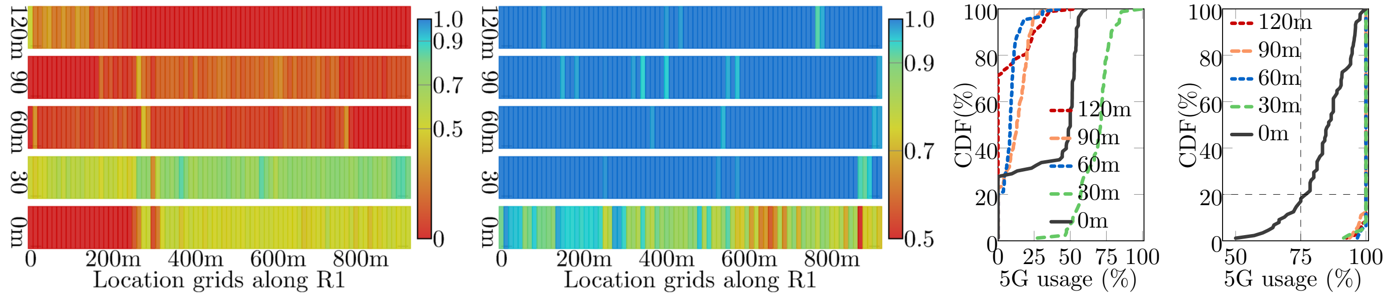

(a) 5G duration ratio map (**A**) (b) 5G duration ratio map (**T**) (c) 5G usage (**A**) (d) 5G usage (**T**)

Fig. 7: RAT-level (5G/4G) usage over R1 (A and T).

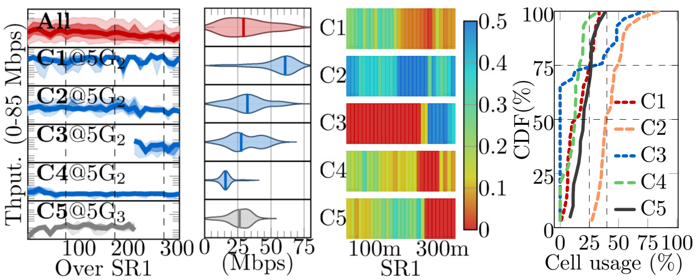

(a) Thput. on SR1 (b) Thput. (c) Duration % (d) CDF

Fig. 8: Performance and usage per cell on subroute SR1 at 120m.

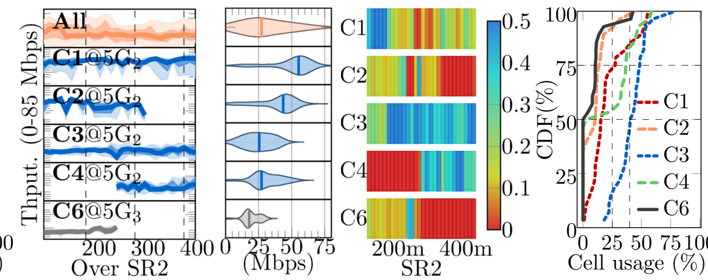

(a) Thput. on SR2 (b) Thput. (c) Duration % (d) CDF

Fig. 9: Performance and usage per cell on subroute SR2 at 90m.

similarly to 4G (about 6 Mbps). It is because 5G performance potentials are not well utilized with current resource allocation on the ground as explained later in §III-E.

We next examine 5G usage by calculating the ratio of 5G duration per location. Fig. 7a and Fig. 7b plot the 5G usage maps of **A** and **T**. Fig. 7c and Fig. 7d show the cumulative distribution functions (CDFs) of 5G usage over all the location grids. For **T**, higher 5G usage is observed at each altitude in the low sky, which contributes better aerial performance than **A**. We see that 5G is used in >95% of time almost everywhere in the sky. On the ground, 5G usage drops below 85%, particularly on the second half of this route R1 (from 500m to 900m). 5G usage is below 75% at 20% of locations on the ground. In contrast, 5G is not often used by **A** in the air, leading to poor aerial performance. Fig. 7c shows low 5G usage except at an altitude of 30 m, where 5G is used at more than half of time at almost all locations. As the altitude goes up above 60 m, 5G is rarely used (below 20%). That explains why the overall performance of **A** above 60 m is significantly worse than at 30 m. In §IV-B, we will analyze how 5G usage at different altitudes varies with respect to RSRP levels of 5G serving cells and 5G configurations. In summary, we see that high performance variance observed in the sky is mainly attributed to the RAT-level diversity for **A**, and its poor 5G usage is likely an dominant contributor to missing high performance potentials offered by 5G in the air (§IV). However, it is not the case for **T**, which uses 5G nearly 100% in the sky. We next show its performance variance is mainly due to the use of distinct cells.

D. Performance Variance Among 5G Cells

We next move one step deeper to the cell level to examine performance variance among different 5G cells. We consider **T** only because inter-cell diversity is less significant for **A** with fewer 5G cells. For **A**, whether 5G is used or not, plays a decisive role to aerial performance, regardless of which 5G cells are used.

[F3] For T, its aerial performance variance is primarily attributed to the use of distinct 5G serving cells.

As shown in Table II, **T** deploys cells over three 5G channels. $5G_2$ is a mid-band channel that uses more frequency resources (bandwidth: 100 MHz) centered on 2600 MHz. $5G_3$ and $5G_4$ are two low-band channels with a much smaller bandwidth (20 MHz). It is not hard to understand that channel bandwidth plays a decisive role on data performance in most instances. 5G cells over $5G_2$ with larger channel bandwidth likely perform better than those over other two 5G channels, which has been validated in our previous work [26].

However, we reveal that **T**'s data performance varies when 5G serving cells change, even though they run on the same channel. Not all serving cells over $5G_2$ yield high data throughput; some cells perform well but others not. To better understand performance variance among cells, we give two illustrative examples over two sub-routes SR1 (Fig. 8) and SR2 (Fig. 9). SR1 is a sub-route of R1 from 20 m to 320 m at an altitude of 120m and SR2 is a sub-route of R1 from 110 m to 410 m at an altitude of 90 m. We consider top-five 5G cells per sub-route. Here, we see four cells (C1 – C4) on $5G_2$ and two cells (C5 and C6) on $5G_3$. Note that four cells@$5G_2$ observed on SR1 and SR2 are the same because these two sub-routes are quite close in the 3D sky. Not all cells are observed anywhere; For example, C2 is seen along SR1 (20m, 320m) at a height of 120 m but partly along SR2 at a height of 90 m; It is primarily impacted by limited radio coverage and partly impacted by HO (unlikely selected at places where its RSRP drops too much).

Fig. 8a and Fig. 9a show the uplink throughput of **T** when the cell is used. Evidently, large performance variance per location (top) is mainly attributed to the use of different serving cells, though cell-level performance does vary. Fig. 8b and

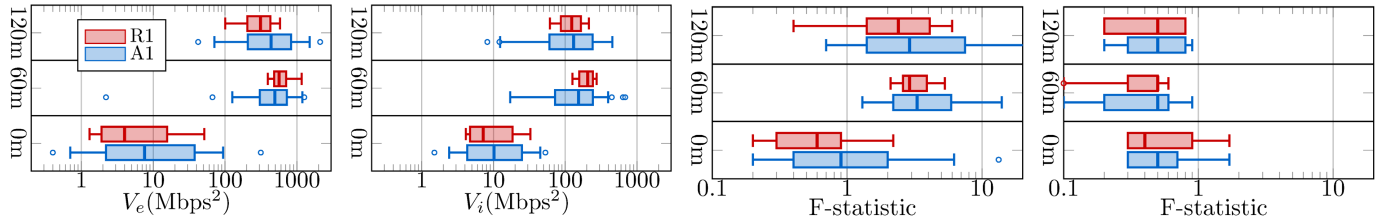

(a) Inter-cell variance (**T**) (b) Intra-cell variance (**T**) (c) F-statistic (**T**) (d) F-statistic (**A**)

Fig. 10: Performance diversity among 5G cells used by A and T.

Fig. 9b show cell-level performance variance is much smaller than location-level one which considers the use of distinct serving cells. Such huge performance diversity among cells is widely observed at other locations. We want to highlight one thing. Despite cell-level performance variance, some cells still statistically outperform others. Generally, C1 performs best with its median throughput above 50 Mbps on both SR1 and SR2; C4 and C6 perform worst with their median throughput below 20 Mbps. It implies that data throughput might quickly lose 30+Mbps when an HO (improperly) switches the serving 5G cell from C1 to C4 (SR1) or C6 (SR2).

As a result, cell selection plays a critical role on the realized performance of **T**; It impacts not only high performance variance but also the likelihood of missing performance potentials. Fig. 8c and Fig. 9c plot the usage ratio of these cells. Surprisingly, we find that the good 5G cells are not selected in most of time. C1 is the best cell but it is not the most popular cell with the highest duration ratio. Conversely, C2 on SR1 and C3 on SR2 are used more often than C1, despite their much lower throughput. As shown in Fig. 8d and Fig. 9d, C1 is used in only <25% of the time on 75% of grids on both SR1 and SR2. However, C2 on SR1 and C3 on SR2 are used in >40% of time on half of grids. These results imply that high performance potential provided by good cell C1 is largely missed due to selection of other cells in practice.

Statistical analysis on performance variance. To quantify throughput variance at all test locations, we apply a one-way analysis of variance (ANOVA) test [27]. ANOVA is a widely used statistical technique for identifying differences among samples (here, throughput samples) from various groups (here, different cells). We apply ANOVA to define three metrics:

Intra-cell variance: $V_i = \sum_{i=1}^{N_c} \sum_{j=1}^{N_i} (s_{ij} - \bar{s}_i)^2 / (N_s - N_c)$,
Inter-cell variance: $V_e = \sum_{i=1}^{N_c} N_i(\bar{t}_i - \bar{t}_{\text{all}})^2 / (N_c - 1)$,
F-statistic: $F = V_e / V_i$.

Here, for a given grid, s_{ij} is the j-th throughput sample of cell C$_i$ and $\bar{s}_i$ is its average throughput. $\bar{t}_{\text{all}}$ is the average throughput of all serving cells on the given grid. N_c is the number of serving cells on the given grid, N_i is the number of throughput samples of cell C$_i$, and N_s is the total number of throughput samples from all serving cells. If inter-cell variance V_e is significantly larger than intra-cell variance V_i, namely, $F \gg 1$, the diversity among different cells is the main contributor of the overall high variance. In this situation cell selection becomes essential to the utilization of performance potentials. Missing good cells results in big data speed drop.

Fig. 10a – 10c show the log-scale boxplots of inter-cell variance V_e, intra-cell variance V_i and F-statistic F observed for **T**. We compare the results over route R1 and across the whole test area A1 (Fig. 2a). We test with two flight altitudes 60 m and 120 m over A1, as well as driving/walking experiments on the ground. First, we observe consistent patterns over R1 and A1. Both inter-cell variance and intra-cell variance are much larger in the sky than on the ground. It is mainly because of slower data speed on the ground. Second, $F > 2$ in the sky and $F < 1$ on the ground, at more than half of test grids over R1 and A1. It means that high aerial performance diversity due to various serving cells is commonly observed throughout the test area. By contrast, Fig. 10d shows that for **A**, performance variance among different 5G cells is less important; $F < 1$ at all altitudes in the sky. To sum up, both operators (**A** and **T**) experience aerial performance variance with distinct causes: for **T**, performance variance mainly comes from the cell-level diversity among different 5G cells, while for **A**, the variance is mainly from the RAT-level distinction between 5G and 4G.

E. Performance Variance at Physical Layer

We finally look into performance variance at the physical layer. Based on 3GPP standard specifications [28], the maximum 5G throughput can be approximately calculated with bandwidth, the number of resource blocks (♯ RBs), and the modulation and coding scheme (MCS), and so on [29]. We thus examine the performance impact of these factors and finally show that cell selection still plays a decisive role to aerial performance variance because physical-layer radio resource allocation is impacted by the serving cells.

[F4] Aerial performance is improved with more physical radio resources in terms of more resource blocks and higher code rates used in the low sky.

We first briefly introduce background on physical radio resource allocation [28]. Allocating physical radio resource needs to determine (1) how many RBs are allocated to each UE, and (2) how many useful bits can be transmitted by a single RB. The former is handled by RB allocation and the latter is handled by MCS adaptation. In 5G/4G networks, RB is the smallest unit of physical radio resources that can be allocated to the UE for data transmission. It is defined in terms of frequency and time, typically consisting of 12 consecutive subcarriers in frequency domain over a certain period of time (usually 0.5ms or 1ms duration in time domain). Increasing the number of RBs enlarges the overall capacity available for data transmission. Given the allocated RB, MCS adaptation tunes the number of useful bits per RB. It involves two key parameters: (a) modulation order, which defines the number of bits transmitted by a single RB, and (b) code rate, which is the ratio between the useful bits and the total transmitted bits with Forward Error Correction (FEC). Evidently, increasing

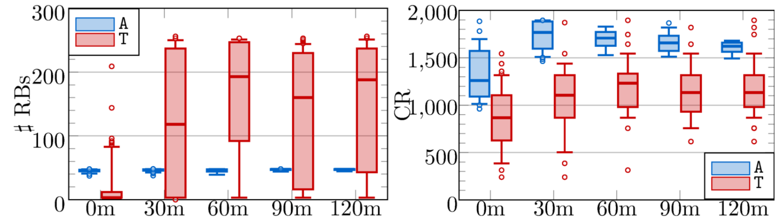

(a) Number of resource blocks (♯ RBs) (b) Coding rate (CR)

Fig. 11: Boxplots of resource block numbers and code rates at all altitudes.

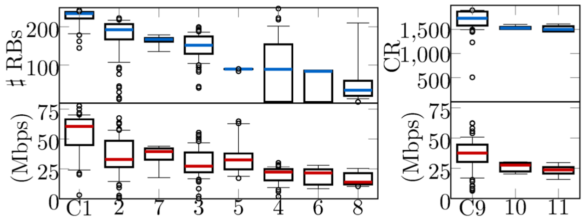

(a) ♯ RBs (120m, **T**) (b) CR (60m, **A**)

Fig. 12: ♯ RBs and CR with different 5G cells.

the modulation order and/or the code rate also contributes to higher throughput.

Fig. 11 shows PHY parameters (♯ RBs and CR) used by **A** and **T** at all altitudes over R1. Evidently, many more RBs are allocated by **T**, while **A** transmits data more efficiently at a higher code rate over a relatively fixed number of RBs. Specifically, we have three observations. First, the number of RBs used by **T** dramatically increases in the air, while the one for **A** remains at the same level. For **T**, the number of allocated RBs is less than 10 on the ground, leading to very low throughput observed in Fig. 6b. In the sky, the median number of RBs increases by more than an order of magnitude, going up to 100 – 200. On the contrary, **A** uses almost the same number of RBs (around 45) at all altitudes, constrained by the narrow channel bandwidth (10 MHz). Second, to compensate limited bandwidth, **A** uses a higher code rate in the sky to enhance transmission efficiency per RB. The code rate grows to 1,500 – 1,800 in the air, which is significantly higher than **T**'s at the same altitudes (900 – 1,300). This explains why **A** is able to achieve performance comparable to **T** over 5G, despite smaller channel bandwidth. Third, in terms of the modulation order, 256QAM is always used in more than 90% of time by both **A** and **T**. As a result, it is not an impact factor on 5G performance, and its results are thus omitted.

We further examine which factor plays a more essential role in impacting aerial performance over 5G.

[F5] For T, performance variance is still dominated by the selection of serving cells because the range of resource block number is impacted by the serving cells.

As shown in Fig. 11a, the number of RBs used by **T** varies a lot up to 200, compared to almost constant use of RBs by **A**. This seems to imply that the diversity in physical resource allocation leads to the huge performance variance observed with **T**. We further perform a breakdown analysis and find such resource allocation diversity is caused by the use of different 5G serving cells. Fig. 12a gives an illustrative instance of the number of RBs and uplink throughput of every 5G serving cell used by **T** at an altitude of 120m over R1. Cells are sorted in the descending order of uplink throughput. Their IDs remain the same as those presented in §III-D. Clearly, the selection of 5G serving cells significantly impacts the allocated number of RBs, thereby resulting in huge performance variance among different cells. Aerial UE can always utilize more than 230 RBs and achieve >60 Mbps when being served by the best cell C1. However, when aerial UE chooses the poor cells (say, C4, C6, and C8), the number of allocated RBs shrinks to 50 – 100 and consequently uplink throughput falls below 25 Mbps.

Compared to **T**, it does not seem necessary for **A** to check its physical parameters as long as 5G is used. As shown in Fig. 11a, the number of allocated RBs remains almost constant (mainly due to the use of this 5G channel). We find that the change in code rates is quite moderate. Fig. 12b shows illustrative results of three 5G cells used by **A** at an altitude of 60 m. Similar results are observed at all other altitudes. First of all, throughput diversity among 5G cells is much smaller than the variance observed for **T**, no more than 10 – 15 Mbps. As a result, for **T**, the selection of 5G serving cell largely impacts the available data transmission resources which further determines performance achieved in the sky. If the UE unfortunately connects to poor cells that provide very limited radio resources, it misses higher performance potentials provided by other cells available.

IV. 5G Performance Potentials Missed in the Sky

In this section, we delve into 5G performance potentials missed in the sky. We first characterize performance potentials missed in the sky, and uncover root causes behind.

A. How Much Performance Potentials are Missed in the Sky?

The challenge to characterize performance potentials missed is that such performance potentials exist but are not utilized in reality, namely, not observed in the measurement experiments. We follow our prior work [24], [25] to run extensive tests at the same locations to learn performance profiles and use the seen to infer the unseen. Specifically, we define good 5G serving cells and assess the missed performance without selecting these good cells when they are present.

[F6] Good 5G cells are missed more often in the sky.

We first determine whether a 5G serving cell is good or not. The rough idea is that the performance of a good 5G cell should be close to the achievable performance potential at the given location. It is hard to obtain the ground truth of such performance potentials on each location. To address it, we use the throughput offered by the best cell on each location to estimate a *lower bound* of the performance potential. For a given cell on a location, we define *ρ-good* to whether the given cell is good or not. Considering intra-cell performance variance, we compare ρ-th percentile throughput of the current cell with $(100-\rho)$-th percentile throughput of the best cell:

$$\rho\text{-}good\ rule:\ T^{\rho}_{current} \geq T^{(100-\rho)}_{best}. \tag{1}$$

If the current cell meets this rule, we call it a ρ-good cell. Clearly, ρ ranges in [50, 100]. The larger ρ, the easier as a good cell. When $\rho = 50$, a 50-good cell is identical to the best cell on the grid. When $\rho = 100$, a 100-good cell with

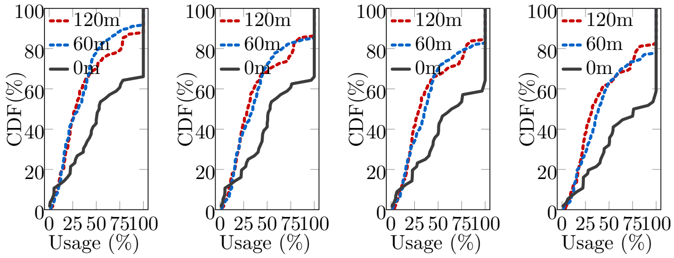

(a) $\rho = 50$ (b) $\rho = 60$ (c) $\rho = 70$ (d) $\rho = 80$

Fig. 13: The usage of good 5G cells over the whole area A1 (`A`).

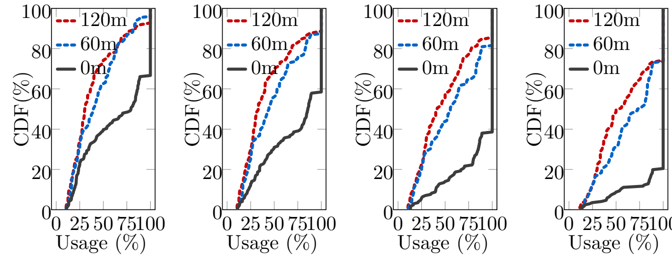

(a) $\rho = 50$ (b) $\rho = 60$ (c) $\rho = 70$ (d) $\rho = 80$

Fig. 14: The usage of good 5G cells over the whole area A1 (`T`).

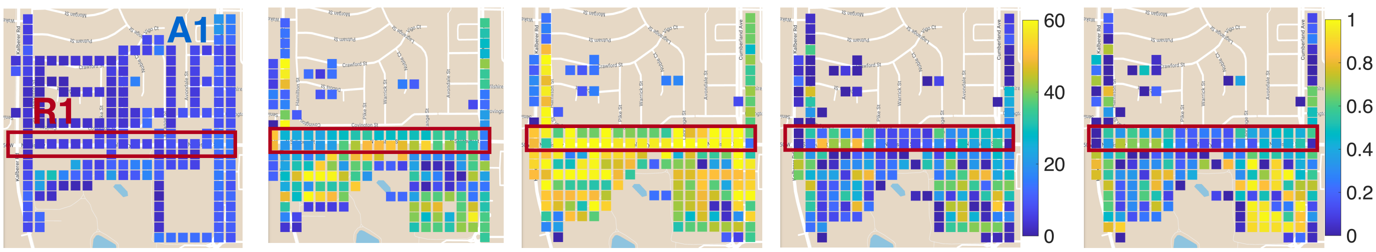

(a) Performance (0 m) (b) Performance (60 m) (c) Potential (60 m) (d) $\Delta^\star$ (60 m) (e) $\delta^\star$ (60 m)

Fig. 15: Realized and missed performance at 60m over the whole area A1 (`T`).

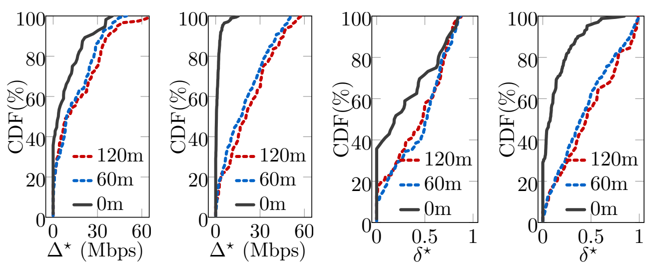

(a) $\Delta^\star$ (`A`) (b) $\Delta^\star$ (`T`) (c) $\delta^\star$ (`A`) (d) $\delta^\star$ (`T`)

Fig. 16: Performance potentials missed over A1.

its maximal throughput larger than the minimal throughput of the best cell is still treated as a good cell.

We next show the usage percentage of good 5G cells using the proposed ρ-good rule. Fig. 13 and Fig. 14 show the results of `A` and `T` in A1, with four ρ values from 50 to 80. For both `A` and `T`, the usage of ρ-good 5G cells in the sky is significantly lower than on the ground, regardless of the ρ value. It indicates that good 5G cells are not frequently used in the sky. Specifically, for `T`, the 80-good cells are used in less than half of time on 30%-50% of aerial locations, while on the ground, the usage is 100% on 80% of locations. Similar results are observed for `A`; Their difference is that the number of 5G cells used by `A` is smaller and most 5G cells are good cells but quite a portion of 5G cells used by `T` are not good. As a result, the value of parameter ρ has a more significant role to the usage of ρ-good cells for `T` than for `A`. When ρ is relaxed from 50 to 80, the usage of good 5G cells by `T` grows at all altitudes, but that of `A` increases very slightly. For instance, at 60m, the median usage of ρ-good cells of `T` rises from 39% ($\rho = 50$) to 73% ($\rho = 80$), but the counterpart of `A` only increases from 31% to 39%. This is because ρ controls which 5G cells are considered as good cells, while the major issue of `A` is that all 5G cells are missed. We set $\rho = 70$ as the default value in the rest of the paper unless specified.

[F7] The missing of good cells prevents more than half of potential throughput from being realized.

We next investigate how much performance potentials are missed due to not selecting good cells. Here, we use `T` as an example to explain how we define realized and missed performance. Fig. 15 shows the maps of realized and missed performance of `T` over the whole testing region A1 at a showcase altitude of 60m. Here we use the median throughput under the cell selection in practice to represent the realized performance on each location. From Fig. 15a to Fig. 15b, 5G does provide much higher throughput on most of locations in the sky where the realized performance is mostly < 20 Mbps on the ground. However, the realized performance is significantly lower than the throughput potential achieved by the best 5G serving cell (Fig. 15c). Aerial UEs could have chance to get additional 40+ Mbps on almost all locations, while the actual performance is usually only 20 – 40 Mbps.

To quantify missed performance potentials, we define two metrics as the upper bound of missed performance potentials:

$$\Delta^\star = T_{best} - T_{worst}, \quad \delta^\star = \Delta^\star / T_{best}. \tag{2}$$

T_{best} and T_{worst} are throughput using the best and worst serving cell on the given grid. $\Delta^\star$ and $\delta^\star$ use the absolute and relative gap to approximate the bound of missed performance. Fig. 15d shows that about 20-40Mbps throughput potentials are missed on half of aerial grids for `T`. It accounts for for more than 50% of the relative miss (Fig. 15e). The results are consistent over R1 and A1. Moreover, the thing is even worse on certain locations of A1. In the bottom right subarea, the relative miss even reaches up to more than 80%. We find such large miss is mainly caused by a configuration problem so that the performance loss is repeatedly and persistently observed in our experiments. More details will be elaborated in §IV-B.

Fig. 16 shows CDFs of the absolute and relative potentials missed at different altitudes of `A` and `T` in A1. We observe that for both `A` and `T`, the impact of missed performance is similar in the sky (namely, at 60 m and 120 m), which is much worse than on the ground. For `A`, at least 20 Mbps throughput is

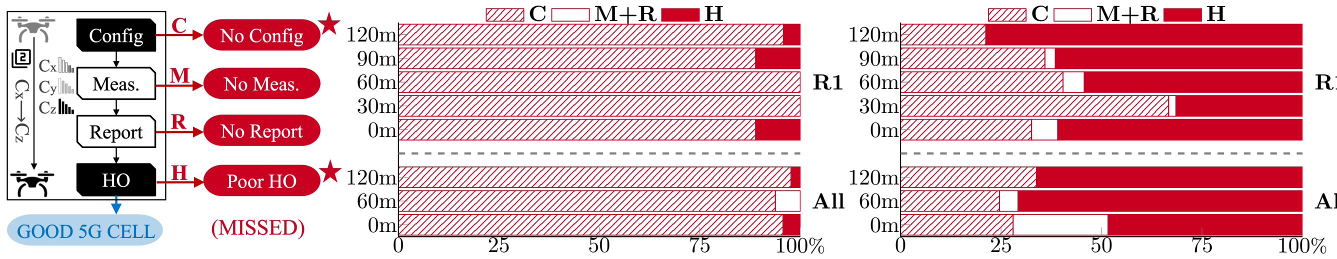

(a) FSM of missing good cells (b) Cause breakdown (**A**) (c) Cause breakdown (**T**)

Fig. 17: Cause analysis of missing good 5G cells and reality check in our study.

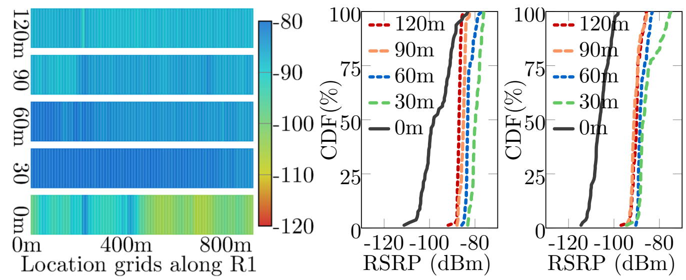

(a) Map of median RSRP (**A**) (b) CDF (**A**) (c) CDF (**T**)

Fig. 18: RSRP of 5G serving cells used by A and T over R1.

missed ($\Delta^{\star}$ >20 Mbps) on more than 35% of locations in the sky, which account for around 50% of performance potential ($\delta^{\star}$ >0.5). The missed performance issue problem in the sky is even more severe for **T**. At least 20 Mbps performance potential is wasted on around half of locations in the sky, and the relative missed performance can even reach up to nearly 100%. By contrast, on the ground, the relative missed performance is less than 25% at 54% and 82% of locations for **A** and **T**, respectively. All these results indicate that the poor cell selection causes much more severe impact on performance for aerial users than ground users for both operators.

B. Why are Good 5G Cells Missed in the Sky?

We next dive into why behind missed performance potentials in the sky. Surprisingly, we find that **A** and **T** suffer from distinct causes though both miss good 5G cells in practice.

We first examine how a HO is performed to figure out where it goes wrong. A standard HO procedure takes four steps: configuration, measurement, report and handover decision/execution. Fig. 17a shows the finite state machine (FSM) for the good and bad branches. Evidently, performance potentials are missed if a HO does not end with selecting a good 5G cell. In another word, this HO instance early exits at the previous C-M-R-H branches. Specifically, there are four branches: (C) good cells are not configured for any measurement; (M) good cells are configured for measurement but they are not measured; (R) good cells are measured but not reported; (H) good cells are reported but not selected as the target cell. We follow a similar approach used by our recent study to analyze performance potentials on the ground [22]. In this study, there are two differences. First, aerial radio channels change and thus the causes for missed performance potentials in the sky vary. Second, more performance potentials are missed in the sky despite common issues for both aerial and terrestrial UE.

[F8] Two operators miss good 5G cells with distinct causes: for A, good 5G cells in the sky are missed mainly due to no configuration, while good cells are missed at either the configuration step or HO decision step for T.

Fig. 17b and Fig. 17c show the breakdown per cause over route R1 and the whole test area A1 for both **A** and **T**. Here we use the default value $\rho = 70$ to determine good cells and not good cells. We use M+R because we cannot tell whether good cells are not measured (M) or measured-but-not-reported (R) if no reporting of good cells is observed. We have two findings.

First, we notice that missing good 5G cells in the sky is unlikely caused by measurement or reporting issues for both **A** and **T**. M+R has a minor impact in all the cases except on the ground in the whole test region. In about 25% of terrestrial instances of **T**, good cells are configured to measure but not reported. It matches with a considerable portion observed in our previous work [22], where RSRP/RSRQ measurements of these good cells are not high enough to trigger the reporting event on the ground. However, it is not the case in the sky because RSRP/RSRQ becomes much higher in the low sky. Fig. 18a shows the median RSRP of 5G serving cells of **A** per grid over R1, and Fig. 18b and Fig. 18c plot the CDFs of RSRP at five altitudes for **A** and **T**. Evidently, the median RSRPs of both operators in the sky are mostly larger than -100 dBm, significantly larger than those on the ground. It is consistent with prior measurement studies on radio quality and this is primarily attributed to line-of-sight propagation and less radio interface [8]. This helps us to understand high 5G usage of **T** in the sky because most 5G candidate cells are qualified with large RSRPs. Poor radio quality on the ground also limits the resource allocation for data transmission, which explains worse 5G performance on the ground (§III-B).

Second, we notice that the primary causes of missing good 5G cells are distinct for **A** and **T**. No 5G configuration is the only dominant reason for **A**. In more than 80% of **A** instances, good 5G cells are missed because there is no valid configuration for any 5G channels, as shown in Fig. 17c. In contrast, poor HO decision is the top-1 reason of missing good 5G cells for **T**. It is responsible for more than half of instances of missing good cells of **T** between 60m and 120m. No 5G configuration is only the secondary cause, which accounts for 20%-40% of instances at each altitude except 30m.

[F9] The why behind no 5G configuration is also different for two operators. For A, no configuration stems from its PCell channel-specific configuration policies, while for T, it is due to its preference for mid-band 5G channels.

We next examine no-configuration instances and we find

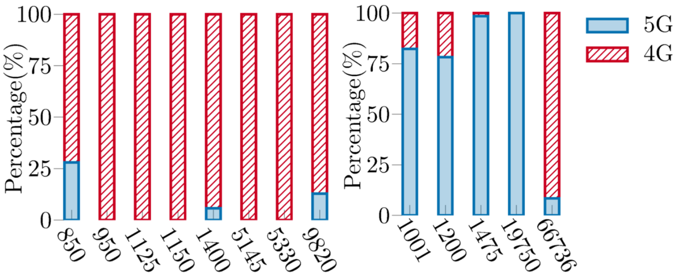

(a) **A** (b) **T**

Fig. 19: Percentage of 5G usage per 4G PCell channel.

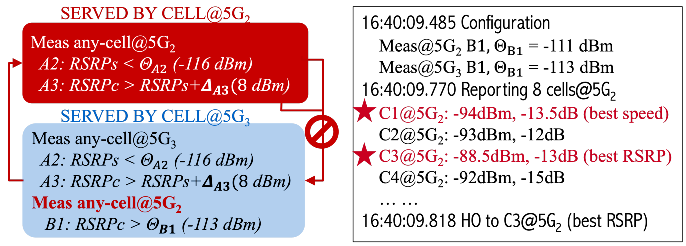

(a) No configuration ($5G_2 \not\to 5G_3$) (b) RSRP-based HO (SR2, 90m)

Fig. 20: Instances of no configuration and poor HO decision (T).

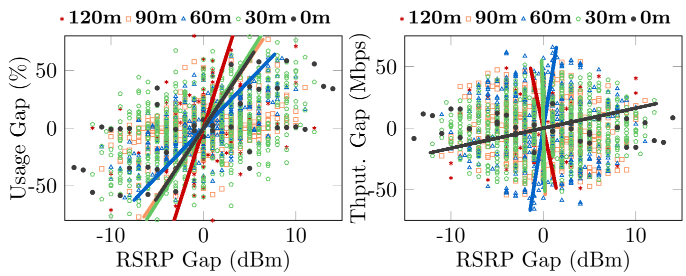

(a) RSRP vs usage (b) RSRP vs throughput

Fig. 21: The relationship between RSRP, usage and throughput at each altitude over R1 (T).

that for **A** and **T**, the why behind no 5G configuration is distinct as well. For **A**, the PCell channel-specific policies result in no 5G configuration. When UE is served by PCell on certain channels, UE will not receive valid configuration of 5G channels, so it cannot measure and use any 5G cells available. In our prior study [22], we first observed this issue in **A**'s 5G networks in Indianapolis and Chicago in 2021. After three years, we found that this issue is still disturbing 5G usage but in another city (West Lafayette, IN). Fig. 19 plots the 5G usage per PCell channel used by **A** and **T**. For **T**, the selection of PCell channel does not restrict the use of 5G, so 5G is still used in more than 75% of time over most of PCell channels. In contrast, for **A**, only one PCell channel (850) is somehow 5G-friendly, which allows 5G usage to exceed 25%. For all other PCell channels, UE either cannot use 5G at all or only use it in less than 10% of time. This indicates that 5G SCell configuration policies used by **A** over multiple 4G PCell channels are likely a nationwide strategy so that the issues are observed in every test city in our prior studies.

The no-configuration instances of **T** are stemmed from its preference on the mid-band channel $5G_2$. Fig. 20a shows configuration items extracted from our traces. They are consistently observed in our study when 5G serving cells are involved. As long as any serving cell@$5G_2$ is used, measurement on channel $5G_3$ will not be configured. However, if any serving cell@$5G_3$ is used, measurement on both $5G_2$ and $5G_3$ will be configured. As a result, no cells@$5G_3$ will be even considered once a cell@$5G_2$ is used, regardless of how well cells@$5G_3$ perform. It indicates that **T** enforces its exclusive preference for mid-band wide channels (here, $5G_2$, 100MHz) over low-band narrow channels (here, $5G_3$ and $5G_4$, 20 MHz). Such preference is not without rationale. Generally, it holds true that good cells@$5G_2$ largely perform better than good cells@$5G_3$ because $5G_2$ uses more bandwidth. However, good cells@$5G_2$ could be missed in cell selection, so they are not always used. The use of one poor cell@$5G_2$ blocks the potential use of good cells@$5G_3$, thereby underutilizing performance potentials deployed in place.

Even worse, the UE quickly loses a good cell@$5G_3$ even though it is used. In Fig. 20a, A2, A3 and B1 are three reporting events regulated by 3GPP [17]. The reporting is trigger if any event is satisfied, for example, if (A2) RSRP of the serving cell is smaller than a threshold θ_{A2} (here, -116 dBm), or (A3) RSRP of a candidate cell is stronger than the RSRP of the serving cell by an offset Δ_{A3} (here, 8 dBm), or (B1) RSRP of a candidate cell is larger than a threshold θ_{B1} (here, -113 dBm). Here, if the UE is served by a cell@$5G_3$, the reporting is triggered with event B1 in presence of any cell@$5G_2$ with RSRP $> \theta_{B1}$ (here, -113 dBm). Given high RSRPs in the sky, it is very easy, if not 100%, to meet the reporting criteria. Note that RSRP thresholds and offsets are configured for the terrestrial cases; Clearly they do not fit well for aerial radio channels with much larger RSRPs. In our experiments, we observe quite a few instances where the UE is initially served by a well-performed cell@$5G_3$, then quickly (within several seconds) switches to a worse cell@$5G_2$ and never gets back to the original good cell. Uplink throughput shrinks by more than 60%, from 50+ Mbps to <20 Mbps. It implies that missed performance potentials are more likely caused by no configuration when cell@$5G_3$ performs better or comparably well.

[F10] For T, good 5G cells are missed largely due to radio-oriented HO decision.

As shown in Fig. 17c, for **T**, poor HO decision (H) is responsible for more than half of HO instances that fail to use good cells despite their presence. It is because HO is mainly based on radio quality measurement but not data performance provided by the cells to be selected. Fig. 20b gives one illustrative instance over a sub-route SR2 at an altitude of 90 m. At the start, it configures to measure cells@$5G_2$ and report RSRP measurements if they are larger than -111 dBm (event B1). After 285 ms, the UE reports 8 5G cells on $5G_2$ including C1@$5G_2$ and C3@$5G_2$. C1@$5G_2$ yields highest data speed (50 – 60 Mbps) but its RSRP is -94 dBm, which is weaker than -88.5 dBm, the RSRP of C3@$5G_2$. Unsurprisingly, C3@$5G_2$ is selected as the new serving cells but C3@$5G_2$ performs worse and thus uplink throughput reduces to 20 – 30 Mbps.

It is not new to blame that radio-centric HO results in missed performance potentials. Our previous studies [24], [25] conducted driving experiments in the same city and showed RSRP-oriented HO fails to select well-performed cells in

4G/4.5G networks. In this work, aerial UE suffers from the same issue but the resulting performance loss is larger due to the change of aerial radio channels. To illustrate this, we define RSRP gap, usage gap, and throughput gap to represent the differences in RSRP (median value), usage, and throughput (median value) between pairwise cells at each location. If usage gap increases as RSRP gap increase, it suggests the preference of selecting cells with higher RSRP; If throughput gap increases as RSRP gap increase, it means selecting cells with higher RSRP is beneficial for achieving better performance. Fig. 21a displays the scatter plot of the RSRP gap and usage gap as well as the linear regression results at each altitude. It clearly shows a positive correlation of RSRP gap and usage gap at all altitudes. This indicates that handover logic is still radio-centric, and cells with higher RSRP are more likely to be selected. However, there is no positive correlation between RSRP and throughput gap in the sky. As shown in Fig. 21b, we can still clearly observe a positive correlation between RSRP and throughput at 0m. However, this relationship becomes less clear or even reverses at altitudes of 30m and 60m, and turns negative at 90m and 120m. These results show that cells with higher RSRP are more unlikely to provide better performance in the air, yet they are still more likely to be chosen due to radio-centric cell selection strategy. In a nutshell, relying on RSRP/RSRQ only is not a wise criterion for selecting well-performed 5G cells in the sky.

V. 5GAIR++: SOLUTION & EVALUATION

Inspired by our previous findings, we propose 5GAIR++, a quick fix solution to tune cell selection for aerial UE to mitigate missed performance over operational 5G networks.

A. The design of 5GAIR++

As described in §IV-B, poor 5G cells are selected mainly because of (1) no 5G configuration and (2) radio-centric HO decision. To tackle these two issues, 5GAIR++ is designed with the following two guidelines: *(1) the configuration step should not exclude candidate cells on certain 5G channels*. All available 5G channels should be configured for measurement and report except that operators intend not to use certain 5G channels. The HO decision step should be responsible for filtering out poor candidate cells. *(2) HO decision should take both radio-related (RSRP/RSRQ) and performance-related metrics (bandwidth and PHY parameters) into consideration* to compare the quality of 5G candidate cells. Fig. 22 depicts the operation flow of 5GAIR++. It incorporates three modules – *aerial 5G profiling*, *altitude-aware configuration* and *good cell prediction* – into the legacy 5G cell selection procedure, ensuring compatibility with standard HO mechanism.

First, 5GAIR++ performs *aerial 5G profiling* to capture essential information of 5G serving cells at each altitude. For each UE, 5GAIR++ monitors its altitude, 5G serving cell, RSRP and performance. Such information is piggybacked in the measurement reports of the serving cells and periodically sent to the base station. 5GAIR++ aggregates the collected data into an offline database, and outputs a cell profiling table. For each observed serving cell at each altitude, 5GAIR++

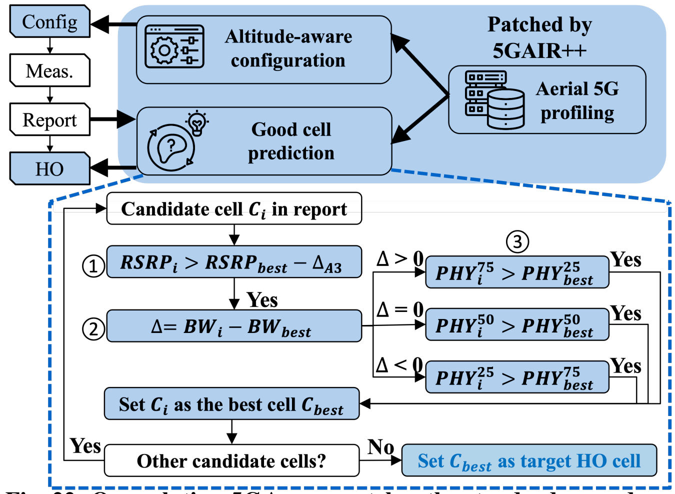

Fig. 22: Our solution 5GAIR++ patches the standard procedure.

leverages historical data to check: (1) the frequency bandwidth of the involved cell, (2) the RSRP range that this cell has been used as a serving cell, and (3) the ρ-th percentile values of resource block number and code rate of this cell in each RSRP range. Here, we consider $\rho = 25$, 50 and 75. To better understand how it works, we give an illustrative example record for cell C1 on channel 5G$_2$ at the altitude of 60 m:

- `Profile`: C1@5G$_2$, altitude = 60m
  - `Bandwidth`: 100 MHz
  - `RSRP range`: [-96dBm, -81dBm]
  - `♯ RBs (25/50/75th-percentile)`:
    * `-100dBm<RSRP≤-90dBm`: 3/202/229
    * `-90dBm<RSRP≤-80dBm`: 169/215/227
  - `CR (25/50/75th-percentile)`:
    * `-100dBm<RSRP≤-90dBm`: 240/1112/1152
    * `-90dBm<RSRP≤-80dBm`: 1177/1402/1708

Next, 5GAIR++ runs *altitude-aware configuration* to improve configuration in the air. At runtime, 5GAIR++ queries the cell profiling table to extract the information of all observed 5G serving cells on the current altitude of UE. All 5G channels observed at the current altitude are configured to prevent the inaccessibility of certain 5G channels due to improper configuration logic. Moreover, 5GAIR++ automatically adjusts the threshold of RSRP based on the altitude to avoid too-low threshold due to terrestrial-based configuration. From the table, we get the list of the lowest RSRP for each 5G serving cell at the current altitude. Accordingly, we set the new RSRP threshold to a reasonable value (e.g., 25-th percentile of all lowest RSRPs) to exclude cells with relatively poor radio strength at this altitude.

In the final HO decision stage, 5GAIR++ performs *good cell prediction*, which takes a new *performance-centric* approach to override the traditional radio-centric cell selection. Specifically, we run a three-level decision tree illustrated in the bottom of Fig. 22. In each round, a candidate cell C_i from the measurement report is compared with a (temporary) best cell C_{best} obtained from previous rounds. The comparison is performed at three levels in turn: RSRP, bandwidth and PHY parameters. On the first level, we examine whether the radio strength $RSRP_i$ of the candidate cell is significantly lower

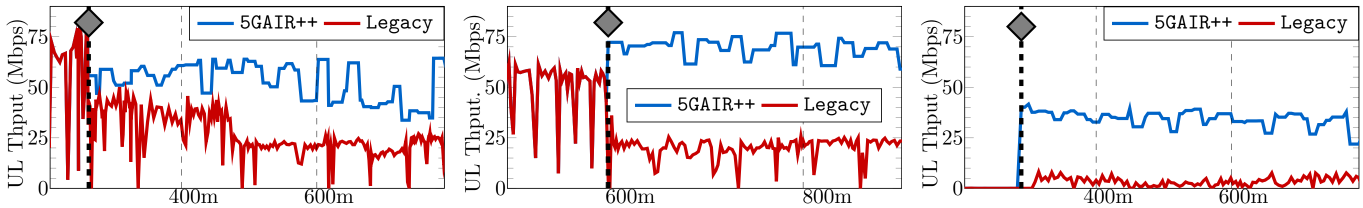

(a) Example 1 (**T**, 90m) (b) Example 2 (**T**, 60m) (c) Example 3 (**A**, 60m)

Fig. 23: Three real-world instances illustrate how 5GAIR++ works and improves uplink throughput (◊: HO changed by 5GAIR++).

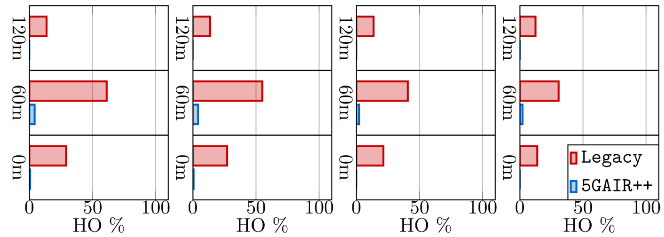

(a) $\rho = 50$ (b) $\rho = 60$ (c) $\rho = 70$ (d) $\rho = 80$

Fig. 24: 5GAIR++ fixes HOs without selecting good cells (A).

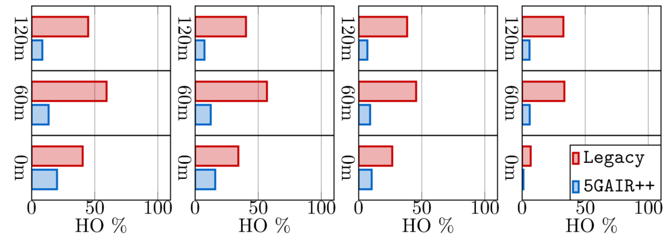

(a) $\rho = 50$ (b) $\rho = 60$ (c) $\rho = 70$ (d) $\rho = 80$

Fig. 25: 5GAIR++ fixes HOs without selecting good cells (T).

than that of the best cell $RSRP_{best}$ by at least Δ_{A3}. Here Δ_{A3} is the offset of A3 event configured by the operator, which has been introduced in §IV-B. This rule not only filters out cells with poor radio strength, but also avoids ping-pong handoffs if 5GAIR++ is not enabled by the target cell. The second level compares the bandwidth of the candidate cell BW_i and the best cell BW_{best}. The bandwidth gap Δ decides which rule is to be used the comparison at the final (third) level. According to the value of Δ, we define three rules:

- Rule I ($\Delta > 0$): $PHY_i^{75} > PHY_{best}^{25}$;
- Rule II ($\Delta = 0$): $PHY_i^{50} > PHY_{best}^{50}$;
- Rule III ($\Delta < 0$): $PHY_i^{25} > PHY_{best}^{75}$.

PHY_i^{ρ} and PHY_{best}^{ρ} are the ρ-th percentile values of the PHY parameter of the candidate cell C_i and the best cell C_{best} in the corresponding RSRP range. Regarding the PHY parameter, we use ♯ RBs for **T** and CR for **A**, based on our analysis in §III-E. With these rules, cells with larger frequency bandwidth are still prioritized for use. However, if these cells cannot provide sufficient radio resources for transmission, they will be replaced by the cells with lower bandwidth but more radio resources. After enumerating all reported cells, 5GAIR++ outputs the best cell as the target cell for handover. By this means, the serving cells selected by 5GAIR++ are more likely to offer good performance for aerial UE.

B. Evaluation with Aerial UE

We take a trace-driven evaluation to assess 5GAIR++, since deploying 5GAIR++ into operational 5G networks is not feasible. Specifically, we take a what-if approach similar to [30]. We run 5GAIR++ on each collected HO instance to determine whether it would recommend an alternative 5G serving cell. If the 5G serving cell changes, we estimate the performance with the new 5G serving cell using its performance profile learned from historical data. We then compare it with the performance using the original serving cells in the legacy HO instance which is obtained in our real-world experiment. We use two applications – bulky file uploading and video live streaming – to assess its potential benefits.

- **Bulky file uploading.** For bulky file uploading, we evaluate the effectiveness of 5GAIR++ using the following metrics: (1) the ratio of HOs without selecting good cells reduced by 5GAIR++, and (2) the absolute and relative throughput gain achieved by 5GAIR++.

[F11] 5GAIR++ effectively fixes more than half of problematic HOs without selecting good cells. It doubles data throughput in 30% – 50% of instances in our study.

We first use three real-world instances to illustrate how 5GAIR++ works in Fig. 23. We compare the uplink throughput observed in the legacy HO instance and the one obtained by 5GAIR++ which fixes the HO without selecting good cells. ◊ marks when 5GAIR++ is triggered to change the HO decision. Each instance is extracted from one run over R1 at different altitudes (60m and 90m) and we show the subroute impacted by 5GAIR++ in Fig. 23. The first two instances (**T**) both start with high data speed but later data speed significantly drops because the legacy HO switches to a poor 5G cell (around 260m in Fig. 23a and around 600m in Fig. 23b. 5GAIR++ recommends other 5G cells which likely perform better than the legacy option based on good cell prediction. Note that in the second instance (Fig. 23b), 5GAIR++ even recommends a new cell which even performs better than the original cell, increasing uplink throughput from below 25Mbps (legacy HO) to 60 – 75 Mbps. The third instance (Fig. 23c) shows how 5GAIR++ helps **A** to fix the legacy HO which attempts to handle a radio access failure. Note that it starts with zero throughput because the legacy operation fails to offer radio access. After a while, the legacy HO is performed to recover data access interruption. It works but offers very low data speed (about 2.8 Mbps, maximum throughout $<$ 8 Mbps) because it uses 4G only. 5GAIR++ recommends using 4G+5G, thereby greatly increasing upload data speed. Remarkably, 5GAIR++ significantly increases the relative data speed, although it cannot increase the absolute data speed a lot for **A**. This matches with our findings that **A** is generally slower than **T** even using 5G. Compared to the gains for **T**, 5GAIR++ cannot make **A** absolutely faster but

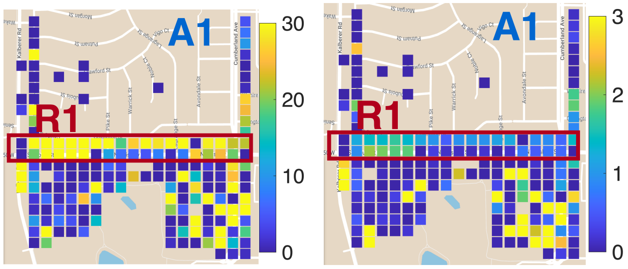

(a) Absolute gain Δ (Mbps) (b) Relative gain γ

Fig. 26: Map of the absolute/relative throughput gains with 5GAIR++ in the showcase at an altitude of 60m (T).

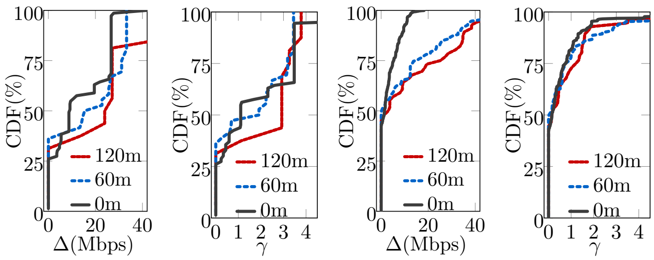

(a) Absolute (**A**) (b) Relative (**A**) (c) Absolute (**T**) (d) Relative (**T**)

Fig. 27: CDF of the absolute and relative gains with 5GAIR++.

yield huge relative gains as the legacy speed is much slower.

We then check how 5GAIR++ can effectively improve the cell selection and reduce the number of HOs without selecting good cells. We show the results using all instances for **A** and **T** in Fig. 24 and Fig. 25. Here, we still test rules with four different $\rho = 50, 60, 70, 80$ used in §IV-A. Clearly, we see that 5GAIR++ is able to effectively reduce problematic HOs for both **A** and **T**. 5GAIR++ is even more effective for **A**; It reduces more than 90% of the legacy HOs without selecting good cells, with the ratio from 13%-61% (legacy) to 0%-4% (5GAIR++). For **T**, 5GAIR++ avoids 80-82% of problematic handovers in the air when $\rho = 70$. Additionally, 5GAIR++ also benefits terrestrial UE, and the probability of selecting poor cells on the ground also declines from 21% to <1% for **A** and from 27% to 10% for **T**.

Finally, we quantify the throughput gain by comparing the median throughput of the new cellset $T_{5\text{GAIR++}}$ and the original cellset T_{legacy} in historical data. We adopt the same metrics in [30]: the absolute gain Δ and the relative gain γ:

$$\Delta = T_{5\text{GAIR++}} - T_{legacy}, \quad \gamma = \Delta / T_{legacy} \tag{3}$$

Fig. 26 uses **T** at an altitude of 60m as the showcase to visualize the throughput gains per location. 5GAIR++ can improve uplink throughput at more than two-thirds of locations at 60m in A1. The absolute throughput gain Δ is higher than 20 Mbps at 30% of locations (Fig. 26a) and the throughput at least doubles ($\gamma > 1$) at about 25% of locations (Fig. 26b). As expected, the performance gain by 5GAIR++ is not confined to specific routes or subareas, but is observed across a broad range of locations. Fig. 27 further shows the results for both **A** and **T** at all three altitudes. We plot the CDF of the absolute and relative gains using all HO instances collected in A1. 5GAIR++ can increase uplink throughput in more than half of legacy HO instances (**A**: 61%-63%, **T**: 50%-54%). In 20%-30% of **T**'s instances in the sky, the throughput is doubled with

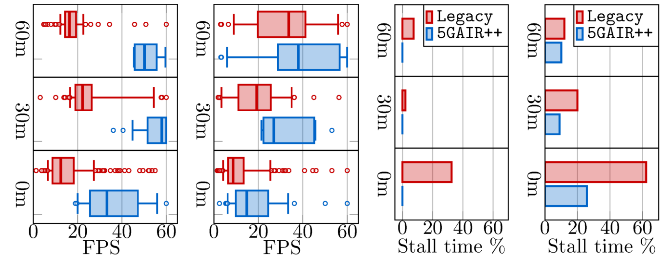

(a) FPS (**A**) (b) FPS (**T**) (c) Stall (**A**) (d) Stall (**T**)

Fig. 28: Performance gains on live streaming with 5GAIR++.

5GAIR++, and throughput gain Δ is higher than 20Mbps. For **A**, despite smaller absolute gains, the relative throughput gains are even larger than 2× in more than half of HO instances and go up one order of magnitude higher.

- **Video live streaming.** We next conduct experiments running 4K video live streaming with constant bitrate (CBR) to evaluate the performance of 5GAIR++. We set up live streaming over WebRTC using open source codes [31]. The test phone attached to the flying drone keeps streaming a 10-second 4K video at 60 fps to our lab server repeatedly. We fly drones over route R2 at three altitudes (0m, 30m, 60m), and record the video frame rate (FPS) and stalls every 100 ms in every run. We take a similar approach to get the KPIs (FPS and stalls) of video streaming achieved by 5GAIR++. For a given legacy HO instance, if the serving cell is changed by 5GAIR++, we use the KPIs observed in other runs which use the new serving cell as the what-if performance.

[F12] 5GAIR++ significantly improves experience of aerial video streaming by increasing the frame rate and reducing the stall time.

Fig. 28 shows how 5GAIR++ improves video streaming KPIs over the legacy solution. In the sky (at 30m and 60m), 5GAIR++ significantly increases the median FPS (**A**: from 16 – 22 fps to 50 – 58 fps, **T**: from 19 – 33 fps to 27 – 38 fps). The improvement is more evident for **A** because its legacy speed is much lower (refer to Fig. 3). Moreover, 5GAIR++ ensures high FPS more reliably for **A** than for **T** (see the boxplots in fig:fps-a and 28b). It matches with our results on bulky file uploading. 5GAIR++ fixes more HOs for **A** and **T** still suffer with low FPS in some instances which are not fixed by 5GAIR++. Meanwhile, 5GAIR++ also greatly reduces video stalls; Stalls are almost negligible for **A** and decrease below 10% for **T**. Here, we want to highlight that 5GAIR++ greatly reduces stalls at the higher FPS, which is enabled by 5GAIR++ and much higher than the original FPS observed in the legacy instances. We also notice that 5GAIR++ improves video streaming on the ground. It is less improved on the ground than in the sky, mainly because the legacy uplink data throughput is much lower on the ground and the improvement is limited though 5GAIR++ boosts uplink throughput on the ground. The maximum throughput on the ground is usually lower than the one observed in the sky in our study.

More remarkably, we notice that **A** benefits more from 5GAIR++ in most cases than **T**. It is because **A**'s performance issues lie in the RAT level, while **T**'s issues are mainly stemmed from the cell-level selection. RAT-level performance

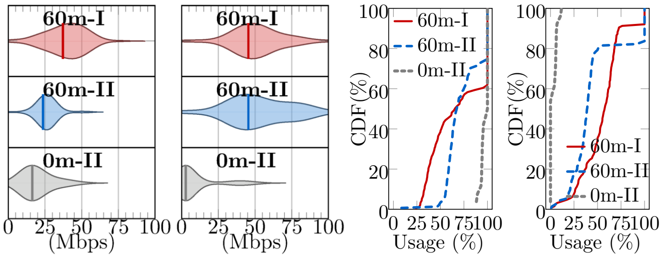

(a) Thput (n2) (b) Thput (n41) (c) Usage (n2) (d) Usage (n41)

Fig. 29: Uplink throughput and band usage in three use scenarios: with one aerial UE only (60m-I), with one aerial UE and one terrestrial UE (60m-II) and with two terrestrial UEs (0m-II).

issues are easier to fix as the prediction is more reliable and fixing no-5G-configuration is simple. 5G consistently outperforms 4G in most of the time, although though comparing performance among cells is much harder, which is impacted by many factors that vary over time and across locations. Nevertheless, 5GAIR++ can easily resolve most problematic HOs through fixing 5G configuration for **A**. For **T**, performance gains mainly count on good cell prediction with historical data for performance profiles. 5GAIR++ shows that it is promising to boost data throughput by selecting better-performed cells with a finer-grained cell-level performance model.

C. Evaluation with Aerial and Terrestrial UE

We next evaluate how 5GAIR++ performs in a new scenario with both aerial and terrestrial UE over operational 5G networks. We aim to understand whether and how co-located active UEs impacts its performance gains.

We notice that for sake of safety, drones in the same airspace should maintain separation distance (250 vertical feet and 2,000 horizontal feet recommended by AOPA [32]). As a result, we do not run experiments with two or more co-located aerial UEs. Instead, we run new experiments with one aerial UE and one terrestrial UE (60m-II), both moving at the same time along the same route (here, `R1`) at different altitudes (60 m and 0 m). Both run the same traffic loads (here, bulky file uploading) but use different remote servers to avoid potential collision on the server-side. For comparison, we also perform experiments with one aerial UE only at 60 m (60m-I) and experiments with two terrestrial UEs (0m-II). We conducted new experiments with **T** in June 2025. We considered **T** only because **T** has better 5G support with more 5G channels and diverse uses (see Fig. 29). We used a new 5G phone model OnePlus 12R as both aerial and terrestrial UE. This new dataset called D3, along with two old datasets, is also released at Github [1].

[F13] The performance of aerial UE is not always impacted by terrestrial UE. It depends on whether they are served by the same 5G cells, which varies with the used frequency channel in operational 5G networks.

Interestingly, we find that the presence of terrestrial UE may not always impact the performance of aerial UE. More specifically, we find that it is because they may not use the same serving cells, which depends on 5G channels/bands used in this study. Fig. 29a and Fig. 29b shows uplink throughput in three use scenarios when the aerial UE is served by 5G

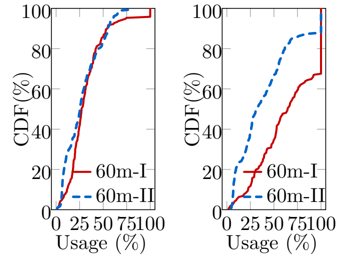

(a) $\rho = 50$ (b) $\rho = 80$

Fig. 30: The usage of good 5G cells with/without ground UE.

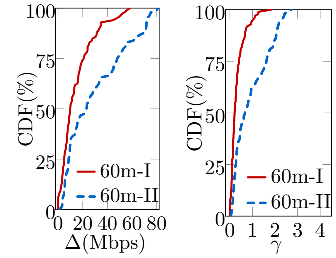

(a) Absolute (b) Relative

Fig. 31: Throughput gains with 5GAIR++ with terrestrial UE.

cells operating over bands n2 and n41. We notice that main 5G bands (and channels) used by **T** changed from our prior measurement in 2024 to our new measurements in 2025. Compared to Table II in dataset D2, **T** uses band n2 (actually, one channel over 1937MHz), not band n71 in dataset D3. We find that the aerial performance indeed degrades in presence of terrestrial UE over band n2. The median uplink throughput drops from around 40 Mbps to 25 Mbps (60m-I vs. 60m-II). But there is no significant performance changes if over band n41 (median throughput around 45 Mbps).

We further figure out why the used 5G bands (actually, channels) make a difference. It turns out that certain serving cells are used by aerial UE (here, n41). Fig. 29c and Fig. 29d plot the usage percentage of serving cells in the sky and on the ground over two bands n2 and n41. We clearly see that terrestrial UE rarely uses 5G serving cells on band n41 but mainly on band n2. On the ground, the usage of serving cells on band n2 exceeds 90% at most of locations, while the one over n41 is nearly zero. In the sky (say, at 60 m), the usage ratio of n2 reduces in the presence of other UE on the ground (60m-I vs. 60m-II). On contrast, the usage ratio of band n41 is not impacted by the presence of terrestrial UE and even slightly goes higher (from 60m-I to 60m-II). As a result, when aerial UE is served by some cells (here, over n41), its perform is not impacted by terrestrial UE. This explains why terrestrial UE only hurts data throughput of aerial UE when its serving cell runs over band n2 in the sky.

[F14] 5GAIR++ yields even higher throughput gains in case the missed performance potential of aerial UE is exacerbated by the presence of terrestrial UE.

We next evaluate the performance gains of 5GAIR++ in case the problem of missing good 5G cells becomes worse due to the presence of terrestrial UE. Here, we examine those cases where 5G serving cells operate over band n2. We do see that the aerial UE misses good serving cels more frequently in presence of other co-located active UE (here, one terrestrial UE). We compare the percentage of selecting good 5G serving cells using two thresholds $\rho = 50$ (Fig. 30a) and $\rho = 80$ (Fig. 30b). From $\rho = 50$ to $\rho = 80$, the presence of one terrestrial UE significantly increases the likelihood of missing good 5G cells (from 50% to 70% at more than half of the locations in Fig. 30b). Surprisingly, we find that the performance gains of 5GAIR++ are larger in presence of this terrestrial UE. In this study, 5GAIR++ improves uplink throughput by over 15 Mbps at 37% of locations, with a gain ratio exceeding 30% without this terrestrial UE (60m-I). In contrast, with this terrestrial UE (60m-II), uplink throughput even improves by more than 30 Mbps at 40% of locations, doubling uplink throughput gains. It is anti-intuitive but not hard to understand. The absolute uplink throughput indeed decreases more in the presence of co-locating active UEs but the missed performance potentials are also larger. As a result, 5GAIR++ enhances data performance more in case more good performance potentials are missed due to the presence of co-locating active UEs.

VI. Related Work

5G/4G measurement for aerial UE. In recent years, quite a number of measurement studies have been conducted to characterize and analyze 5G/4G connectivity and performance for drones ([6]–[15], [33]–[36]). Earliest efforts were traced back to 3GPP's work item in 2017 [37], which aimed to understand potentials and issues of supporting drones over 4G and resulted in TR36.777 [8], the first 3GPP technical report over field trials performed by major stakeholders (Qualcomm, Nokia, Ericsson, and Huawei, etc). These early field trials were mostly performed at a single site (operated by these companies) and focused on charactering radio quality in the low sky. They were followed by many field tests centered on radio quality, interference and even channel propagation models ([9], [10], [33]–[36]). Recent studies have shifted their focus to measure data performance over operational cellular networks: 4G [8]–[10], [12], [13] and 5G [6], [7], [13]–[15]. However, they simply reported the absolute throughput observed in their measurement studies, which vary drastically due to distinct network deployment and environmental factors. In contrast, we not only characterize performance observed but also analyze performance unobserved, namely, performance potentials missed in the low sky. Moreover, we analyze why and propose a solution to enhancing 5G aerial performance.

Performance potentials missed for terrestrial UE. Missed performance potentials were first revealed in our measurement study over 4G [25], followed up by several recent studies [22], [24], [30], [38]. All these studies show that cell selection should take the blame for performance potentials missed for terrestrial UE. In addition, several studies have measured and analyzed the practice of cell selection for terrestrial UE [23], [39]–[41]. Our work is inspired by these efforts but targets at aerial UE, which experiences distinct radio channels in the low sky and misses more potentials as cell selection is not properly configured for aerial radio channels.

Performance potentials missed for aerial UE. Our preliminary work [26] is the first and only study on missed performance potentials for aerial UE, to our best knowledge. This work has substantially extended our prior efforts with new results from two US operators and compared the impacts of operator-specific configuration and operations. Moreover, we have developed an enhanced solution inspired by new insights.

VII. Conclusion and Discussion

In this work, we present our efforts to measure, analyze and enhance aerial performance for drones that fly in the low sky over operational 5G networks. Through extensive measurements with two US operators (AT&T and T-Mobile), we have several interesting findings on 5G performance potentials realized and missed in the sky. On one hand, we observe larger uplink data speed in the sky which turns higher aerial performance potentials into reality; On the other hand, we also notice huge data performance variance which implies that performance potentials are likely missed in practice. Huge performance variance is primarily stemmed from the use of various serving cells, which is attributed to much higher RSRP in the low sky than on the ground. However, stronger radio coverage turns out into a double-sided sword because well-performed cells might not be always selected for use. Despite operator-specific issues, current practice in 5G networks are designated for terrestrial use and do not well work for aerial UE. We devise 5GAIR++, a quick patch to fix HO configuration and decision. We validate that it is promising to pursue more aerial performance potentials.

There are some remaining issues. First of all, this work focuses on data performance of a single aerial UE. We do not consider heavy aerial traffic with many UEs, which is not allowed by FAA and AOPA now. However, in the foreseen future with many more drones in the sky, heavy traffic from many drones in the same airspace will make this problem of missed performance potentials much more complex. Load balancing among multiple candidate cells will change the premise that missing good cells in the legacy cell selection is mainly caused by improper radio signal quality comparison. Our prior study with 4.5G on the ground [24] shows that the performance gains may decrease when the number of active UE keep increasing. Despite the distinction in the sky, it likely holds true with many more aerial UEs in the sky. Second, though we spent huge efforts in measuring aerial 5G performance in the sky, the scale is still limited. It is because running aerial experiments are much harder and more time-consuming, constrained by very limited drone battery lifetime (up to 30 minutes) and line-of-sight flight requirement. We are developing a crowd-sourcing platform which will facilitate other groups in the research community to use the same/similar methods to accelerate data collection across various environments and build up a more comprehensive understanding of 5G performance in the sky.

Acknowledgments. We thank all anonymous reviewers for their constructive feedback. The work was partially supported by NSF grants CNS-1750953 and CNS-2112471.

References

[1] "5G Measurement Datasets: Aerial 5G in West Lafayette," https://github.com/mssn/5GMeas-Dataset, 2025.

[2] GSMA, "Mobile-Enabled Unmanned Aircraft," https://www.gsma.com/iot/wp-content/uploads/2018/02/Mobile-Enabled-Unmanned-Aircraft-web.pdf, Feb 2018.

[3] D. Mishra and E. Natalizio, "A Survey on Cellular-Connected UAVs: Design Challenges, Enabling 5G/B5G Innovations, and Experimental Advancements," *Computer Networks*, vol. 182, 2020.

[4] G. Geraci, A. Garcia-Rodriguez, M. M. Azari, A. Lozano, M. Mezzavilla, S. Chatzinotas, Y. Chen, S. Rangan, and M. Di Renzo, "What will the Future of UAV Cellular Communications be? A Flight from 5G to 6G," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 3, pp. 1304–1335, 2022.

[5] FAA, "PART 107: Small Unmanned Aircraft Systems," https://www.ecfr.gov/current/title-14/chapter-I/subchapter-F/part-107.

[6] R. Muzaffar, C. Raffelsberger, A. Fakhreddine, J. L. Luque, D. Emini, and C. Bettstetter, "First Experiments with a 5G-Connected Drone," in *DroneNet*, 2020.
[7] S. Horsmanheimo, L. Tuomimäki, V. Semkin, S. Mehnert, T. Chen, M. Ojennus, and L. Nykänen, "5G Communication QoS Measurements for Smart City UAV Services," in *16th European Conference on Antennas and Propagation (EuCAP)*, 2022.
[8] 3GPP, "TR36.777: Enhanced LTE Support for Aerial Vehicles," Jan. 2018, (Release 15).
[9] S. D. Muruganathan, X. Lin, H.-L. Maattanen, Z. Zou, W. A. Hapsari, and S. Yasukawa, "An Overview of 3GPP Release-15 Study on Enhanced LTE Support for Connected Drones," *arXiv preprint arXiv:1805.00826*, 2018.
[10] S. D. Muruganathan, X. Lin, H.-L. Määttänen, J. Sedin, Z. Zou, W. A. Hapsari, and S. Yasukawa, "An Overview of 3GPP Release-15 Study on Enhanced LTE Support for Connected Drones," *IEEE Communications Standards Magazine*, vol. 5, no. 4, pp. 140–146, 2021.
[11] S. Hayat, C. Bettstetter, R. Muzaffar, and D. Emini, "An Experimental Evaluation of LTE-A Throughput for Drones," in *DroneNet*, 2019.
[12] M. Gharib, S. Nandadapu, and F. Afghah, "An Exhaustive Study of Using Commercial LTE Network for UAV Communication in Rural Areas," in *ICC*, 2021.
[13] S. Homayouni, M. Paier, C. Benischek, G. Pernjak, M. Reichelt, and C. Fuchsjäger, "Field Trials and Design Insights of Cellular-Connected Drones," in *VTC2021-Fall*, 2021.
[14] A. Festag, S. Udupa, L. Garcia, R. Wellens, M. Hecht, and P. Ulfig, "End-to-End Performance Measurements of Drone Communications in 5G Cellular Networks," in *VTC2021-Fall*, 2021.
[15] M. Gharib, B. Hopkins, J. Murrin, A. Koka, and F. Afghah, "5G Wings: Investigating 5G-Connected Drones Performance in Non-Urban Areas," in *PIMRC*, 2023.
[16] 3GPP, "TS22.125: Unmanned Aerial System (UAS) support in 3GPP," December 2023, v19.1.0.
[17] ——, "TS38.331: NR; Radio Resource Control (RRC)," March 2023, v16.12.0.
[18] ——, "TS36.331: E-UTRA; Radio Resource Control (RRC)," March 2023, v16.12.0.
[19] ——, "TS37.340: NR; Multi-connectivity; Overall description; Stage-2," July 2022, v16.10.0.
[20] ——, "TS23.501: System Architecture for the 5G System," 2023, v16.16.0.
[21] ——, "Carrier aggregation on mobile networks," https://www.3gpp.org/technologies/carrier-aggregation-on-mobile-networks, August 2022.
[22] Y. Liu and C. Peng, "A Close Look at 5G in the Wild: Unrealized Potentials and Implications," in *INFOCOM*, 2023.
[23] A. Hassan, A. Narayanan, A. Zhang, W. Ye, R. Zhu, S. Jin, J. Carpenter, Z. M. Mao, F. Qian, and Z.-L. Zhang, "Vivisecting mobility management in 5g cellular networks," in *Proceedings of the ACM SIGCOMM 2022 Conference*, 2022, pp. 86–100.
[24] H. Deng, Q. Li, J. Huang, and C. Peng, "iCellSpeed: Increasing Cellular Data Speed with Device-Assisted Cell Selection," in *MobiCom*, 2020.
[25] H. Deng, K. Ling, J. Guo, and C. Peng, "Unveiling the Missed 4.5G Performance In the Wild," in *HotMobile*, March 2020.
[26] Y. Liu, J. Huang, and C. Peng, "The Sky is Not the Limit: Unveiling Operational 5G Potentials in the Sky," in *IWQoS*, 2024.
[27] E. R. Girden, *ANOVA: Repeated measures*. sage, 1992.
[28] 3GPP, "TS38.306: NR; User Equipment (UE) radio access capabilities," April 2024, v17.8.0.
[29] "5G NR Throughput calculator," https://5g-tools.com/5g-nr-throughput-calculator/.
[30] Q. Li and C. Peng, "Reconfiguring Cell Selection in 4G/5G Networks," in *ICNP*, 2021.
[31] "WebRTC-Streamer," https://github.com/mpromonet/webrtc-streamer.
[32] Aircraft Owners and Pilots Association (AOPA, "Safe separation defined," https://www.aopa.org/news-and-media/all-news/2017/may/15/safe-separation-defined, 2017.
[33] X. Lin, V. Yajnanarayana, S. D. Muruganathan, S. Gao, H. Asplund, H.-L. Maattanen, M. Bergstrom, S. Euler, and Y.-P. E. Wang, "The Sky is not the Limit: LTE for Unmanned Aerial Vehicles," *IEEE Communications Magazine*, vol. 56, no. 4, pp. 204–210, 2018.
[34] X. Lin, R. Wiren, S. Euler, A. Sadam, H.-L. Maattanen, S. Muruganathan, S. Gao, Y.-P. E. Wang, J. Kauppi, Z. Zou *et al.*, "Mobile network-connected drones: Field trials, simulations, and design insights," *IEEE Vehicular Technology Magazine*, vol. 14, no. 3, pp. 115–125, 2019.
[35] A. Fjodorov, A. Masood, M. M. Alam, and S. Pärand, "5G Testbed Implementation and Measurement Campaign for Ground and Aerial Coverage," in *18th Biennial Baltic Electronics Conference (BEC)*, 2022.
[36] J. Luo, P. Zhao, F.-C. Zheng, and L. Li, "Delay Evaluation for Cellular-Connected Drones: Experiments and Analysis," in *VTC2022-Fall*, 2022.
[37] 3GPP, "Work Item "Study on Enhanced Support for Aerial Vehicles"," June 2017, 3GPP TSG RAN Meeting #76, RP-171050.
[38] E. Coronado, S. Siddiqui, and R. Riggio, "Roadrunner: O-ran-based cell selection in beyond 5g networks," in *2022 IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2022, pp. 1–7.
[39] Z. Zhang, Y. Liu, Q. Li, Z. Liu, C. Peng, and S. Lu, "Dependent Misconfigurations in 5G/4.5G Radio Resource Control," in *CoNext*, 2023.
[40] Q. Li, Z. Zhang, Y. Liu, Z. Tan, C. Peng, and S. Lu, "CA++: Enhancing Carrier Aggregation Beyond 5G," in *MobiCom*, 2023.
[41] Y. Liu, G. Guo, and C. Peng, "Demystifying Secondary Radio Access Failures in 5G," in *HotMobile*, 2024.

Yanbing Liu (S'23) received his M.S. and B.S. degree in Department of Electronic Engineering and Information Science at University of Science and Technology of China. He is currently pursuing Ph.D. degree in Department of Computer Science at Purdue University. His research interests are in the area of mobile networking, with a focus on 5G/6G networks measurement and design.

Jingqi Huang received her B.S. degree in Internet of Things Engineering from Beijing University of Posts and Telecommunications in 2018, followed by an M.S. degree in Electrical and Computer Engineering from the University of California, San Diego. She received her Ph.D. degree in Computer Science from Purdue University. Her research interests lie in mobile networking, with a particular focus on the measurement and analysis of 5G networks.

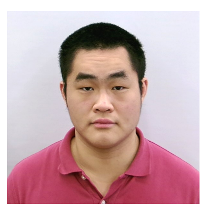

Ziyu Li received his B.S. degrees in Computer Science and in Data Science from Purdue University in 2025. His research interests include mobile networking, system security, and AI systems, with recent work focusing on 5G network measurement.

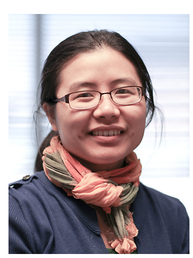

Chunyi Peng (S'09 – M'13 – SM'17) received her Ph.D. at Department of Computer Science, University of California, Los Angeles in 2013. She is currently an Associate Professor at Department of Computer Science, Purdue University, West Lafayette, IN, USA. Prior to that, she worked as an Assistant Professor at Department of Computer Science and Engineering at the Ohio State University and an associate Researcher at Microsoft Research Asia. Her research interests are in the broad areas of mobile networking, system and security, with a recent focus on renovating 5G access technologies, AI for networks, mobile network security, mobile edge computing mainly for autonomous drones and robots.