The Sky is Not the Limit:
Unveiling Operational 5G Potentials in the Sky

Yanbing Liu Jingqi Huang Chunyi Peng
Department of Computer Science, Purdue University, West Lafayette, IN, USA
{liu3098,huan1504,chunyi}@purdue.edu

***Abstract*—In this work, we present our measurement study to characterize and analyze operational 5G performance potentials for cellular-connected drones that fly in the low sky. We not only measure aerial performance observed over an operational 5G network (here, T-Mobile, one major 5G operator in the US), but also quantitively assess potentials missed in the sky. Different from prior measurement studies, we compare 5G performance potentials realized and missed in the low sky and on the ground. 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 can have been even much faster but such potentials are not fully utilized in the sky). We delve into root causes behind missed potentials and find that current 5G cell selection is designed for terrestrial scenarios and misses good candidate cells under aerial radio channel conditions. We thus devise a patch solution called 5GAIR to pursue more 5G potentials in the low sky and validate its effectiveness over real-world traces (released at [1]).**

I. INTRODUCTION

Cellular-connected drones are gaining momentum with their emerging and thrilling uses such as aerial surveillance, traffic monitoring, site inspection, post-disaster rescue, transport and logistics, and to name many [2]. Drones need to transfer a variety of data like videos, images, sensor data, commands and application-specific results 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 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 (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, [6], [7] observed the peak rate of 600 – 700 Mbps when the drone flied around one 5G cell tower, but [14], [15] saw only several tens of Mbps over 5G; [7] found that HOs significantly impacted data performance but [14] believed not; [14] reported that data speed slightly dropped at higher altitude but [15] observed the opposite trend (up to 400 ft). 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 environment, which do vary a lot across distinct field trials.

However, it raises an important problem in understanding aerial performance over operational cellular networks, and gaining insights to enhance 5G support for drones in the sky. *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? Why if not? How to further enhance performance over the existing networks?*

In this work, we attempt to answer the above questions through our measurement study over one operational 5G network in the US. We focus on uplink data throughput because most drone applications (say, aerial surveillance) demand for high uplink data speed [16]. Table I summarizes our main findings. Different from previous measurement studies, we not only quantitively characterize aerial performance observed in our study but also investigate performance potentials which are not observed but missed in the wild (§III). 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 that can be applied to understand aerial 5G performance elsewhere. Moreover, we reason about performance potentials which are available but missed in reality, by comparing performance observed at runtime (namely, performance potentials realized) with reference performance profiles. Interestingly, we find that performance variance, instead of the absolute performance value, is a stronger indicator of performance potential utilization. We delve into an in-depth cause analysis and identify critical factors impacting aerial 5G performance and potentials missed in operational 5G networks. Inspired by our findings, we propose 5GAIR, a quick patch to enhance 5G performance for drones in the low sky (§IV). Our evaluation over real world traces validates that 5GAIR can effectively mitigate the missed performance issue and double uplink throughput in 25% of aerial instances.

Release. Our codes and datasets are released at [1].

II. BACKGROUND

We first introduce necessary background on cellular radio access and then present 5G features observed in this study.

	Category	Description	Figure(s)
F1 **F6**	**What** (§III-A)	5G is much faster in the sky than on the ground but its data throughput varies wildly. 5G good cells performs better in the sky; But they are not always selected for use and thus missing 5G performance potentials is not rare.	Fig. 3 Fig. 4, 5
F3 **F4** **F5**	**Why** (§III-B, III-C)	Aerial performance variance is primarily contributed by the use of distinct 5G serving cells. Good 5G cells are missed more often in the sky. The missing of good cells causes more than half of the potential throughput being unrealized.	Fig. 7, 8, 9 Fig. 10, 11 Fig. 12, 13
F6 **F7**	**Fix** (§IV)	Current practice fails to select good cells mainly due to no configuration and poor HO decision. 5GAIR fixes >50% of problematic handovers and doubles throughput in 25% of instances.	Fig. 14b, 16 Fig. 19, 20

TABLE I: Summary of our main findings (marked as F1 – F7).

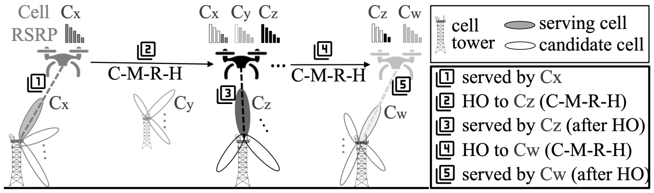

Fig. 1: Cellular-connected drones get "anytime, anywhere" connectivity in the low sky over the existing 5G/4G networks.

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 through standard handover (HO) procedures [17], [18], which is the same for both terrestrial UE and aerial UE. Basically, each HO procedure relies on radio quality measurements (say, RSRP or RSRQ) to determine whether to re-select a new serving cell. It typically takes four steps: (C) configuration, (M) measurement, (R) reporting, (H) handover decision and execution. At the start, the UE is served by one cell, which sends configured parameters to customize the subsequent measurement and reporting steps including but not limited what cells to be measured as well as the criteria to report measurements. These criteria are defined as reporting events (say, A1-A6, B1, B2) by comparing the measured RSRP/RSRQ of the serving cell and candidate cells [17], [18]. Afterwards, the UE measure cells nearby on configured channels and report their measurements when the criteria are met. Finally, the serving cell decides whether to execute a HO over all the reported measurements and switches to another if applicable. Fig. 1 shows two HO instances ($C_x \rightarrow C_z$, $C_z \rightarrow C_w$) where the RSRP of the serving cells drops but other candidate cells offer better radio quality.

5G features observed in this study. 5G follows the above common procedure to establish and mitigate radio access while adopting several advanced features observed in this study. First, 5G networks use both 5G and 4G (two RATs) to serve the UE over a technique called dual connectivity [19]. Moreover, we observe that US operators run 5G primarily in Non-Standalone (NSA) mode, where 4G acts as the master RAT and 5G offers secondary radio access [20]. Second, each RAT (5G or 4G) uses carrier aggregation to allow more than one serving cells [21]. As a result, single UE is served by a set of serving cells, not one cell. In this study, a serving cellset consists of two cell groups (4G+5G) if 5G is used, otherwise one cell group over 4G (4G only). Each cell group consists of one primary cell (PCell) and several secondary cells (SCells). 4G PCell is responsible for configuring and performing HOs. If 4G PCell switches to a new one, other cells are added later.

III. WHAT DOES AERIAL 5G PERFORMANCE LOOK LIKE?

In this section, we present aerial 5G performance observed in our extensive measurements over T-Mobile, one major 5G operator in the US, and then analyze why 5G performs so.

Methodology and dataset. We run measurement experiments primarily in one 1 Km × 1 Km area in a US city (city name is hidden for anonymity) (Map in Fig. 2b). This test area is a typical residential zone with single family houses (top), parks and sport fields (center) and apartments/condos (bottom). It is fully covered by T-Mobile 5G, which operates over low-band and mid-band frequencies (see Table III). Since commodity drones with 5G connectivity are not available yet, we use a drone (here, DJI Phantom 4 Pro) carrying a 5G phone (here, Google Pixel 5) to measure aerial performance (Fig. 2a). We fly drones at different altitudes up to 120m (below 400 ft allowed by FAA [5]). We test with various flight routes including two selected routes R1 and R2, as well as many random routes in the low sky over the test area. We also run driving experiments (only along the roads) and use the same phone to measure performance on the ground (at an altitude of 0 m). In this study, we are interested in uplink data speed because many drone applications need to transfer heavy traffic (e.g., videos, images and sensor data). In all the experiments, the test phone repeatedly uploads bulky files (50MB each) to our lab server and measures uplink speed.

We have conducted experiments sporadically from Nov 2023 to Feb 2024 and collected data logs over 27.8 hours (Table II). In total, we have collected 25K HO instances with more than 1M throughput samples and 5.4M RSRP/RSRQ samples. We see 954 cells including 159 cells over 5G and 795 cells over 4G. It indicates dense cell deployment with many candidate cells (several tens of cells or more) at each location. It matches with recent measurement results [22], [23].

Duration	Distance	♯ Cell	♯ HO	♯ Thput	♯ RSRP/Q
27.8 hr	206 Km	954 (5G:159)	25,092	1.0 M	5.4 M

TABLE II: Dataset statistics.

A. 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

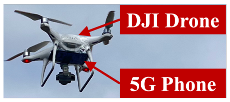

(a) Test equipment

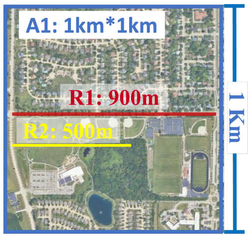

(b) Map

Fig. 2: Exp. settings.

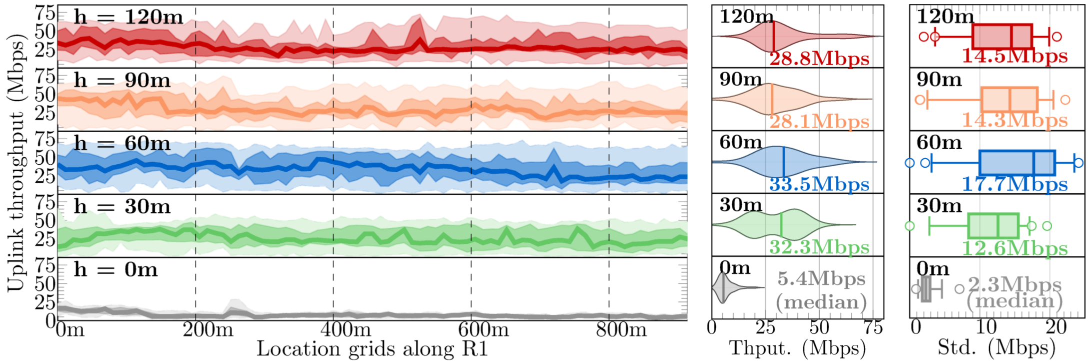

(a) 5/25/50/75/95th-percentile throughput per location (b) Over all grids (c) Std. per grid

Fig. 3: Uplink data throughput observed over R1 at different heights (h = 0m, 30m, 60m, 90m, and 120m).

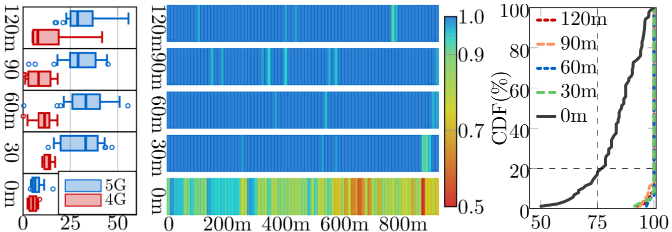

(a) Thput. (b) Map of 5G usage (c) CDF

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

	Band	Ch. Freq	Ch. BW	♯ SCells	Usage
5G$_1$	mid-band, n41	2600 MHz	100 MHz	21	72.9%
5G$_2$	low-band, n71	626 MHz	20 MHz	8	23.2%
5G$_3$	low-band, n71	649 MHz	20 MHz	1	3.9%

TABLE III: Information of three 5G channels observed over R1.

route R1. All main findings are also consistently observed in the test area (§III-B, §III-C). 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). Fig. 3 shows uplink throughput per location (grid size: 10 m).

[F1] 5G is much faster in the sky than on the ground but its data throughput varies wildly.

We plot the 5/25/50/75/95th-percentile of uplink throughput per location in Fig. 3a and the distribution of all throughput samples at all the locations in Fig. 3b. Evidently, 5G is much faster in the sky than on the ground at the same location. Here, aerial UE achieves 20 – 45 Mbps in the sky but the same phone gets <10 Mbps on the ground. The median throughput grows by 5.2× from 5.4 Mbps to 28.1 Mbps (or more).

Interestingly, we see that data throughput varies wildly in the sky. Aerial performance fluctuates at runtime from several Mbps (say, 5th-percentile) to several tens of Mbps (say, 95th-percentile). It is more evident in Fig. 3c. We calculate the standard deviation (std) of uplink throughput per location and shows its boxplot in Fig. 3c. Here, larger performance variance (12.6 – 17.7 Mbps) is observed in the sky, almost an order of magnitude higher than the terrestrial one (2.3 Mbps).

There are two implications on 5G performance potentials. On one hand, higher performance potentials are realized in the sky. On the other hand, considerable performance potentials might be missed in reality as 5G networks offer low data speed at locations where faster speed is available. We next delve into why and reveal 5G potentials realized and missed in practice.

[F2] 5G good cells are used more often, which contributes to better aerial performance. However, current practice cannot always select such well-performed 5G cells and thus missing 5G performance potentials is not rare.

Selecting good serving cells contributes to better performance in the sky. In particular, 5G performance potentials are realized at two levels: RAT and cell.

∘ *At the RAT level, 5G is faster and used more often in the sky than on the ground.* Fig. 4a compares data throughput over 5G and 4G. Note that 5G is never used alone because T-Mobile runs NSA 5G in our test area. All the performance over 5G is provided by a set of serving cells over 4G and 5G (4G+5G). We compare it with data performance provided by 4G-only cellsets. Unsurprisingly, 5G (more precisely, 4G+5G) is much faster than 4G only. In all the aerial scenarios, 5G performs significantly better than 4G only, boosting the median throughput by 2.4× – 3.7× from 5–15 Mbps to 28.9–33.9 Mbps. It indicates 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. It is because 5G performance potentials are not well utilized on the ground and we will explain it later in §IV-A.

We further examine 5G usage by calculating the ratio of 5G duration per location. Fig. 4b plots the map of 5G usage and Fig. 4c shows the cumulative distribution functions (CDFs) of 5G usage over all the location grids. Clearly, better performance on the sky is attributed to higher 5G usage. We see that 5G is used at more than 95% of time almost everywhere in the sky. However on the ground, 5G usage drops below 85%, particularly on the second half of this route R1 (from 500m to 900m). Fig. 4c shows that 5G usage is below 75% at 20% of locations on the ground.

∘ *At the cell level, 5G cells with larger frequency bandwidth (that likely perform better) are more often used in the sky.*

We see that T-Mobile deploys cells over three 5G channels in this study (Table III). 5G$_1$ is a mid-band channel that uses more frequency resources (bandwidth: 100 MHz) centered on 2600 MHz. 5G$_2$ and 5G$_3$ are two low-band channels with

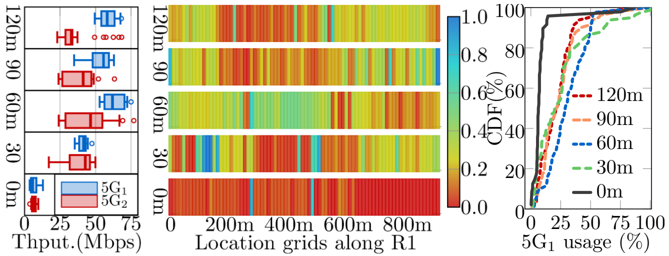

(a) Thput. (b) Usage of good cells over $5G_1$ (c) CDF

Fig. 5: Data throughput and usage of good cells on $5G_1$ over R1.

a much smaller bandwidth (20 MHz). There is no surprise that cells over $5G_1$ largely perform better than those over other two 5G channels (see an illustrative example in Fig. 8a). It is not hard to understand that frequency bandwidth plays a decisive role on data performance in most instances. We do observe exceptions because actual performance is also impacted by many factors like radio quality, runtime traffic loads, radio resource allocation, and so on. They change over time, resulting in dynamic performance. We will later show that data performance is impacted more by the serving 5G cell (as well as its channel bandwidth). Over route R1, we see 30 serving cells ($5G_1$: 21, $5G_2$: 8 and $5G_3$: 1) over three 5G channels (Table III). $5G_1$ is indeed used more, at 72.9% of time in all the experiments at all five altitudes.

Interestingly, not all serving cells over $5G_1$ yield high data throughput. Some cells perform well but others not. Given performance diversity (§III-B), we use good cells out of many cells on the same channel to show performance potentials achieved in reality. Good cells are defined in §III-C. Here, we want to highlight that the good cells using more bandwidth generally perform better. Fig. 5a compares data throughput of good cells on two channels $5G_1$ and $5G_2$; $5G_3$ is ignored because it is rarely used (only one serving cell). We have two observations. First, throughput gains are more evident at higher altitudes. Good cells over $5G_1$ outperforms those over $5G_2$ in the sky and the gain decreases as the drone descends to the ground. There is no evident gain at 30 m (though $5G_1$ slightly performs better in the worst case) and there is almost no difference on the ground. It implies that performance potentials contributed by larger bandwidth are not well utilized on the ground or in the very low sky. Second, good cells over $5G_1$ are not very often used in practice. Fig. 5b and Fig. 5c show the ratio of using good cells over $5G_1$ at various locations. Though they are used more often at higher altitude, the actual ratio is below 50% at most places. Essentially, good cells are indeed available in place but they are not just selected for use due to current practice in 5G networks (§IV-A). By comparing Fig. 4b and Fig. 5b, we get a rough feel on the likelihood of using 5G but not using good cells over 5G. Surprisingly, aerial UE is served by sub-optimal cells at most time even though aerial UE gets better performance than terrestrial UE.

B. Performance Diversity Among 5G Cells

Before we delve into 5G performance potentials missed in the sky, we reveal huge diversity of performance among cells.

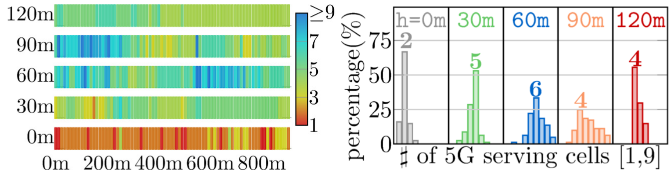

(a) Map of 5G serving cell number (b) Ratio of each cell #

Fig. 6: Number of 5G serving cells observed over R1.

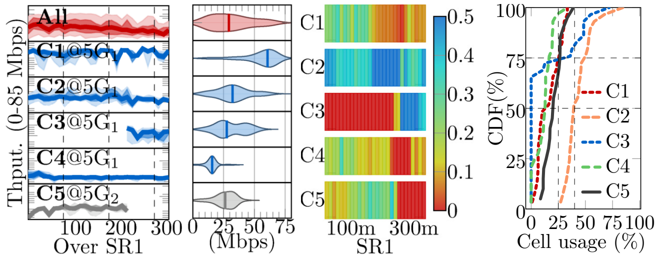

(a) Thput. on SR1 (b) Thput. (c) Usage (d) CDF

Fig. 7: Performance and usage per cell on subroute SR1 at 120m.

We first show that there are more 5G cells available for cell selection in the sky. Fig. 6a plots the number of unique 5G serving cells observed on each location over R1. Fig. 6b shows the percentage of location grids with the number of 5G serving cells varying from 1 to ≥9. Evidently, there are more serving cells in the sky. More than four 5G cells are observed to serve aerial UE at more than half of grids at each altitude. For all altitudes in the sky, at least four serving cells are available on >80% of grids. In contrast, there are no more than two serving cells observed on 70% of grids on the ground. Later we will explain that it is because 5G cells offer higher RSRP in the sky (Fig. 15 in §IV-A).

[F3] Aerial performance variance is primarily contributed by the use of distinct 5G serving cells.

Data performance varies when 5G serving cells change, even though they run on the same channel. To better understand performance diversity per cell, we give two illustrative examples over two sub-routes SR1 (Fig. 7) and SR2 (Fig. 8). 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 $5G_1$ and two cells (C5 and C6) on $5G_2$. Note that four cells@$5G_1$ 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. 7a and Fig. 8a show uplink throughput 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. 7b and Fig. 8b 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

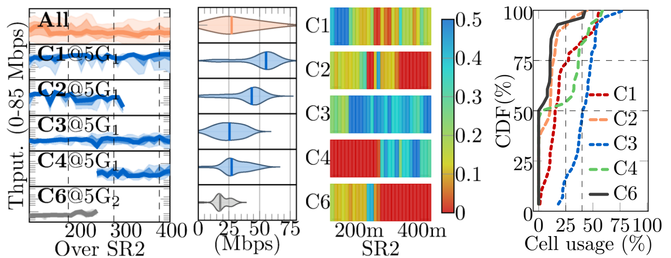

(a) Thput. on SR2 (b) Thput. (c) Usage (d) CDF

Fig. 8: Performance and usage per cell on subroute SR2 at 90m.

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; It impacts not only high performance variance but also the likelihood of missing performance potentials. Fig. 7c and Fig. 8c plot the usage ratio of these cells. Surprinsingly, 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. 7d and Fig. 8d, 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 more than 40% of duration on half of grids. These results imply that the high performance potential provided by the good cell C1 is largely missed due to selection of other cells in practice.

To quantify throughput diversity at all test locations, we apply a one-way analysis of variance (ANOVA) test [24]. ANOVA is a widely used statistical technique for identifying differences among samples (here, throughput samples) from various groups (here, different cells). Here, 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. 9 shows the log-scale boxplot of inter-cell variance V_e, intra-cell variance V_i and F-statistic F at all the grids over R1, R2 and A1. As shown in the map (Fig. 2b), R2 is another shorter route (500 m) and A1 is the whole test area. We test with two flight altitudes 60 m and 120 m over R2 and A1, as well as driving/walking experiments on the ground. First, we observe consistent patterns over R1, R2 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 50% of test grids over R1, R2 and A1. It means that high aerial performance diversity due to various serving cells is commonly observed throughout the test region.

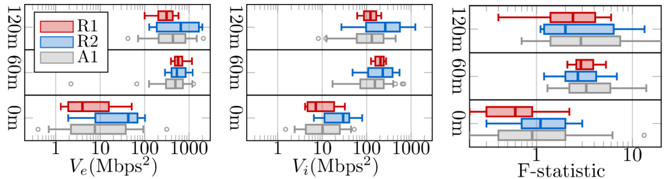

(a) Inter-cell variance (b) Intra-cell variance (c) F-statistic

Fig. 9: Performance diversity of 5G cells over R1, R2, and A1.

C. *Performance Potentials Missed in Practice*

We delve into 5G performance potentials missed in practice. The challenge to characterize performance potentials missed is that such performance potentials exist but not utilized in reality, namely, not observed in the measurement experiments. We follow [25] to run extensive tests to learn performance profiles and use the seen to infer the unseen. In particular, we define good 5G serving cells and assess the missed performance without selecting good cells when they are present.

[F4] 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, if not impossible, 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:

$$\rho\text{-}good\ rule:\ T_{current}^{\rho} \geq T_{best}^{(100-\rho)}. \quad (1)$$

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 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. 10 and Fig. 11 show the results over R1 and A1, with four ρ values from 50 to 80. The results over R1 and A1 are similar. When ρ is relaxed from 50 to 80, the usage of good 5G cells significantly increases at each altitude, especially on the ground. On the ground, 50-good 5G cells are always used (usage = 100%) on around half of

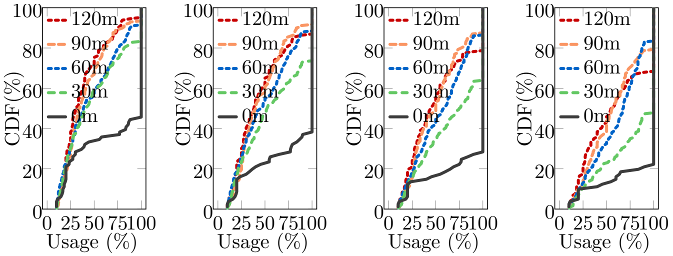

(a) $\rho = 50$ (b) $\rho = 60$ (c) $\rho = 70$ (d) $\rho = 80$

Fig. 10: The usage of good 5G cells in R1.

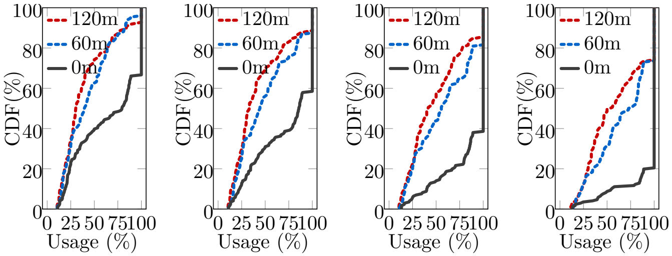

(a) $\rho = 50$ (b) $\rho = 60$ (c) $\rho = 70$ (d) $\rho = 80$

Fig. 11: The usage of good 5G cells in A1.

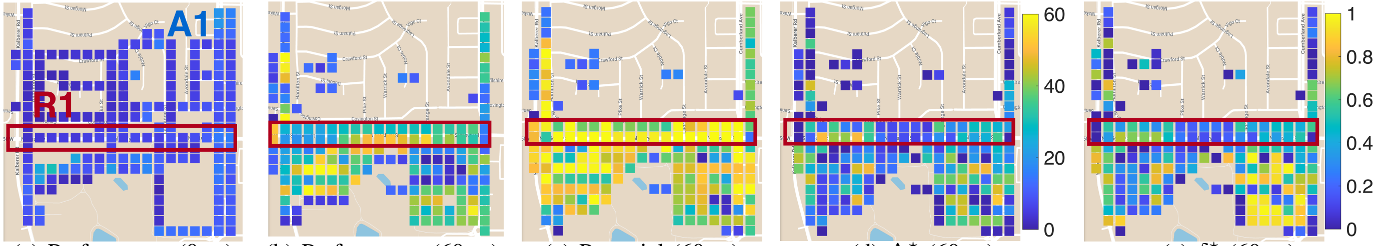

(a) Performance (0 m) (b) Performance (60 m) (c) Potential (60 m) (d) $\Delta^\star$ (60 m) (e) $\delta^\star$ (60 m)

Fig. 12: Realized and missed performance at 60m in A1.

locations. When ρ rises to 80, the usage of 80-good 5G cells is 100% on 80% of locations. It is not hard to understand that data performance among various cells on the ground are quite close; More serving cells, if not all, are treated as good cells as ρ becomes larger. We notice that 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. Good cells are used at less than 50% of time on more than half of aerial location grids when $\rho = 50$ or 60. When ρ grows to 70 and 80, the good 5G cells are still missed in at least 25% of time on more than 60% of locations (except 80-good cells at 30m). We set $\rho = 70$ as the default value in the rest of the paper unless specified.

[F5] The missing of good cells causes more than half of the potential throughput being unrealized.

We next investigate how much performance potentials are missed due to not selecting good cells. Fig. 12 shows the maps of realized and missed performance 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. 12a to Fig. 12b, 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. 12c). 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 data 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.

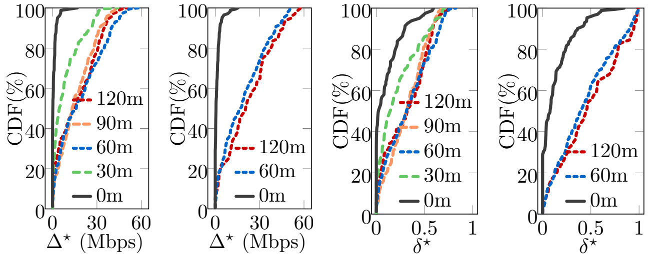

(a) $\Delta^\star$ (R1) (b) $\Delta^\star$ (A1) (c) $\delta^\star$ (R1) (d) $\delta^\star$ (A1)

Fig. 13: Missed performance on different altitudes in R1 and A1.

Fig. 12d shows that about 20 – 40Mbps throughput potentials are missed on around half of aerial grids. It accounts for for more than 50% of the relative miss (Fig. 12e). 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-A.

Fig. 13 shows CDFs of the absolute and relative potentials missed at different altitudes. We observe that the impact of missed performance is similar above 60 m (namely, at 60 m, 90 m and 120 m). It is worse than in the lower sky (say, 30 m) and on the ground. On R1, at least 20 Mbps throughput is missed ($\Delta^\star$ >20 Mbps) on more than 40% of locations at 60-120m, which account for around 40% of performance potential ($\delta^\star$ >0.4). Up to 75% of throughput potential is not realized on certain locations in the sky. By contrast, the relative missed performance is less than 25% on 85% of locations on the ground. Compared with R1, The missed performance problem in the sky of A1 is even more severe. As shown in Fig. 13d, at least 50% of performance potential is wasted on around half of locations in the sky. The relative missed performance can even reach up to nearly 100% in A1. All these results indicate that the poor cell selection causes much more severe impact

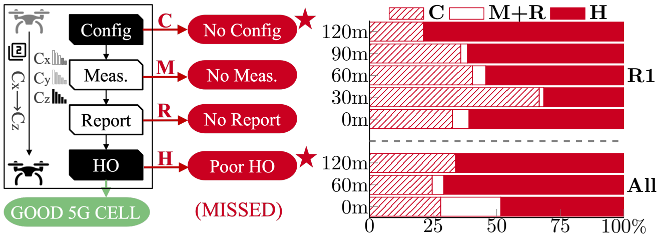

(a) FSM of Missing Good Cells (b) Breakdown

Fig. 14: Cause analysis and reality check in our study.

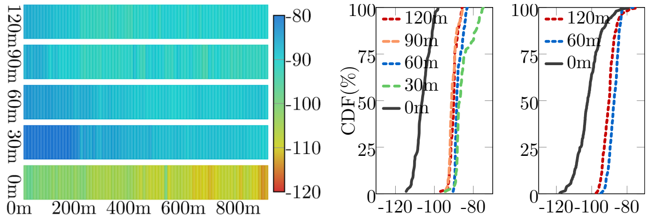

(a) Map of median RSRP (R1) (b) CDF (R1) (c) CDF (All)

Fig. 15: RSRP of 5G serving cells at all test locations.

on performance for aerial users than ground users.

IV. HOW TO ENHANCE AERIAL 5G PERFORMANCE?

In this section, we dive into why behind missed performance potentials in the sky. We find that good cells are missed primarily at two steps: configuration and handover decision. We further devise 5GAIR, a quick patch to tune cell selection for aerial UE over operational 5G networks.

A. *Why are Good Cells Missed in the Sky?*

We 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. 14a 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]. It is easy to understand because the handover procedure is the same for both aerial and terrestrial UE. In this study, there are two main 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.

[F6] The utilization of performance potentials depends on the selection of 5G serving cells. Good cells in the sky are missed primarily at configuration and HO decision steps.

Fig. 14b shows the breakdown per cause over route R1 and the whole test area. 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 three findings.

First, we notice that missing good 5G cells in the sky is unlikely caused by measurement or reporting issues, which is different from the results on the ground. M+R has a minor impact in all the cases except on the ground in the whole test region. In about 25% of instances, 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. 15a shows the median RSRP of 5G serving cells per grid over R1 and Fig. 15b plots the CDFs at five altitudes. Evidently, the median RSRPs 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 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-A). Good cells are missed in the sky primarily due to no configuration or poor handover decision.

Second, we examine no-configuration instances and find that all are stemmed from T-Mobile's preference on $5G_1$. Fig. 16a 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@$5G_1$ is used, measurement on channel $5G_2$ will not be configured. However, if any serving cell@$5G_2$ is used, measurement on both $5G_1$ and $5G_2$ will be configured. As a result, no cells@$5G_2$ will be even considered once a cell@$5G_1$ is used, regardless of how well cells@$5G_2$ perform. It indicates that T-Mobile enforces its exclusive preference for mid-band wide channels (here, $5G_1$, 100MHz) over low-band narrow channels (here, $5G_2$ and $5G_3$, 20 MHz). Such preference is not without rationale. Generally, it holds true that good cells@$5G_1$ largely perform better than good cells@$5G_2$ because $5G_1$ uses more bandwidth. However, good cells@$5G_1$ are not always used as shown in Fig. 5. The use of one poor cell@$5G_1$ blocks the potential use of good cells@$5G_2$, thereby underutilizing performance potentials deployed in place.

Even worse, the UE quickly loses a good cell@$5G_2$ even though it is used. In Fig. 16a, 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@$5G_2$, the reporting is triggered with event B1 in presence of any

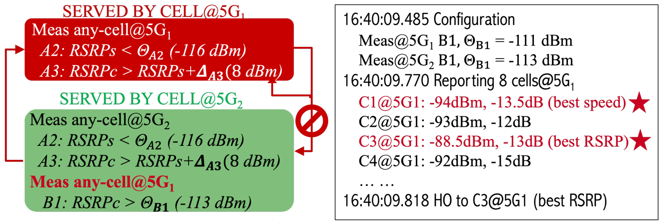

```
16:40:09.485 Configuration
  Meas@5G1 B1, ΘB1 = -111 dBm
  Meas@5G2 B1, ΘB1 = -113 dBm
16:40:09.770 Reporting 8 cells@5G1
  C1@5G1: -94dBm, -13.5dB (best speed)★
  C2@5G1: -93dBm, -12dB
  C3@5G1: -88.5dBm, -13dB (best RSRP)★
  C4@5G1: -92dBm, -15dB
  … …
16:40:09.818 HO to C3@5G1 (best RSRP)
```

(a) No configuration ($5G_1 \not\to 5G_2$) (b) RSRP-based HO (SR2, 90m)

Fig. 16: Instances for no configuration and poor HO decision.

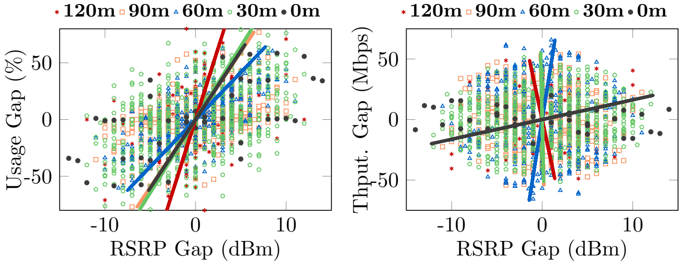

(a) RSRP vs usage (b) RSRP vs throughput

Fig. 17: The relationship between RSRP, usage and throughput at each altitude in R1.

cell@$5G_1$ 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@$5G_2$, then quickly (within several seconds) switches to a worse cell@$5G_1$ 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@$5G_2$ performs better or comparably well. This explains why the peak ratio goes up to 67% (at 30m over R1) because good cells over these two channels perform similarly (Fig. 5a).

Third, good cells are still missed even though they are measured and reported. Poor HO decision (H) is responsible for more than half of 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. 16b shows one illustrative instance observed at a sub-route SR2 at 90 m. At the start, it configures to measure cells@$5G_1$ and report RSRP measurements if they are larger than -111 dBm (event B1). After 285 ms, the UE reports 8 5G cells on $5G_1$ including C1@$5G_1$ and C3@$5G_1$. C1@$5G_1$ yields highest data speed (50 – 60 Mbps) but its RSRP is -94 dBm, which is weaker than -88.5 dBm, the RSRP of C3@$5G_1$. There is no surprise that C3@$5G_1$ is selected as the new serving cells but C3@$5G_1$ performs worse and thus uplink throughput reduce to 20 – 30 Mbps.

It is not new to blame that radio-centric HO results in missed performance potentials. Our previous studies [25], [26] 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. 17a 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. 17b, 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.

B. 5GAIR*: Solution & Evaluation*

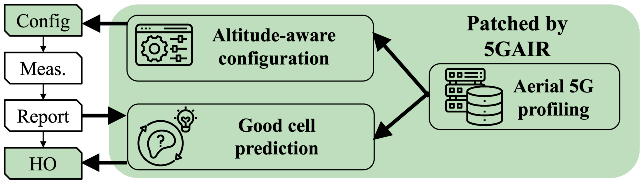

Fig. 18: Overview of our solution 5GAIR.

To mitigate missed performance in the sky, we propose 5GAIR, a quick fix solution to solve the issue of no-configuration and poor HO decision. 5GAIR incorporates three patches (*aerial 5G profiling*, *altitude-aware configuration* and *good cell prediction*) into the legacy 5G cell selection procedure, which ensures the compatibility with standard HO mechanism. Fig. 18 illustrates the operation flow of 5GAIR.

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 uses its historical data to check: (1) the RSRP range that this cell has been used as serving cell, and (2) the duration ratio of this cell serving as a good 5G cell in each RSRP range. We give an illustrative example record for cell C1 on channel $5G_1$ at altitude of 60m:

- C1@$5G_1$, altitude 60m
  - `RSRP range`: [-96dBm, -81dBm]
  - `Good cell probability`:

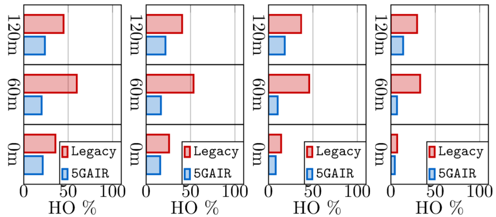

(a) $\rho = 50$ (b) $\rho = 60$ (c) $\rho = 70$ (d) $\rho = 80$

Fig. 19: HO ratio without selecting good cells.

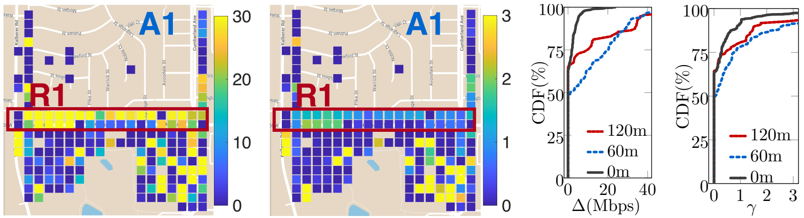

(a) Absolute (60m) (b) Relative (60m) (c) Absolute (d) Relative

Fig. 20: Absolute gain Δ and relative gain γ with 5GAIR in A1.

- `-100dBm<RSRP≤-90dBm`: 54%
- `-90dBm<RSRP≤-80dBm`: 82%

Next, with this table as input, 5GAIR implements *altitude-aware configuration* to improve configuration in the air. At runtime, 5GAIR queries the table to extract the information (channel, RSRP range and good cell probability) of all observed 5G serving cells on the current altitude of UE. All 5G channels observed on the current altitude are configured to prevent the inaccessibility of certain channels (e.g. C2 in subarea A2) 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 adopts a new performance-centric approach based on *good cell prediction* to replace the traditional radio-centric cell selection. When UE prepares a measurement report, 5GAIR gets the altitude as well as the measured RSRP value for each reported cell. Using these information as input, 5GAIR checks the corresponding entry of the cell profiling table to predict the probability to be a good serving cell for each reported cell. For example, if cell C1 is reported with -92dBm RSRP at 60m, 5GAIR will use the 82% from the table as the prediction value of the good cell probability if this cell is selected in this HO. Then, 5GAIR sorts all the reported cells and selects the cell with the highest probability in prediction. As the result, the serving cells selected by 5GAIR are more likely to provide good performance for aerial users.

Evaluation. Since deploying 5GAIR in operational 5G networks is not feasible, we adopt a trace-driven evaluation to assess its potential benefits. We run 5GAIR on each collected handover instance to determine whether it would recommend an alternate 5G serving cell. If the 5G serving cell changes, we estimate the performance with the new 5G serving cell using the historical data. We then compare it with the performance of the original serving cell to evaluate the performance gain.

[F7] 5GAIR effectively fixes more than half of problematic HOs without selecting good cells. It doubles data throughput in 25% of instances in our study.

Fig. 19 shows the ratio of handovers selecting not good cells with 5GAIR and legacy cell selection in A1. Here we still test rules with four different $\rho = 50, 60, 70, 80$ used in §III-C. At altitude of 60m, the ratio of handovers selecting not good cells is reduced by more than two-thirds by 5GAIR, from 33%-60% to 7%-20%. At 120m, 5GAIR is slightly less effective but still can avoid around half of "not-OK" handovers. Additionally, 5GAIR doesn't hurt the cell selection of ground users, and the probability of selecting not good cells on the ground also slightly declines from 7%-36% to 3%-21% with 5GAIR.

To quantify the throughput gain, we compare 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 [27], absolute gain ($\Delta = T_{5\text{GAIR}} - T_{legacy}$) and relative gain ($\gamma = \Delta / T_{legacy}$), to evaluate the performance gain by 5GAIR. Fig. 20a and Fig. 20b use 60m as the showcase altitude to visualize the throughput gains per location. 5GAIR can benefit the throughput on more than two-thirds of locations at 60m in A1. Fig. 20b shows that on 25% of locations, 5GAIR can at least double the throughput ($\gamma > 1$). The absolute throughput gain Δ is higher than 20 Mbps in 30% of locations (Fig. 20a). Moreover, the performance gain by 5GAIR is not confined to specific routes or subareas, but is observed across a broad range of locations. We further extend the evaluation to all altitudes. Fig. 20c and Fig. 20d show the CDF of Δ and γ for handover instances at each altitude in A1. 5GAIR impacts the throughput in 36%-50% of handover instances. In 20%-25% of instances in the sky, the throughput is doubled with 5GAIR, and throughput gain Δ is higher than 20Mbps.

V. RELATED WORK

5G/4G measurement for aerial UE. In recent years, a number of measurement studies have been conducted to characterize and analyze 5G/4G connectivity and performance for drones ([6]–[15], [28]–[33]). Earliest efforts were traced back to 3GPP's work item in 2017 [34], 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 industry. These early field trials were mostly performed at a single site 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], [28]–[33]). 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, which vary drastically due to distinct network deployment and environmental factors. In contrast, our measurement study not only characterizes performance observed but also analyzes performance unobserved, namely, performance potentials missed in the low sky. Moreover, we 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 [26], followed up by several recent studies [22], [25], [27], [35]. 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], [36], [36]–[38]. 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.

VI. CONCLUSION

In this paper, we present the first measurement study to reveal and characterize 5G performance potentials realized and missed in the sky. Different from conventional terrestrial scenarios, serving cells work at much higher RSRP in the low sky due to distinct aerial radio channels. However, it turns out into a double-sided sword. On one hand, we do observe higher data performance in the sky which turns higher aerial performance potentials into reality; On the other hand, we also notice huge data performance variance due to the use of various serving cells which do not always perform well. Current practice in 5G networks are designated for terrestrial use and do not well work for aerial UE. We devise a quick patch to fix configuration and HO decision and validate that it is promising to pursue more aerial performance potentials.

Acknowledgements. We thank all anonymous reviewers for their constructive feedback. The work has been 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, 2024.

[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, S. Jin, A. Narayanan, R. Zhu, A. Zhang, W. Ye, J. Carpenter, Z. M. Mao, Z.-L. Zhang, and F. Qian, "Vivisecting mobility management in 5g cellular networks," in *SIGCOMM'22*, 2022.

[24] E. R. Girden, *ANOVA: Repeated measures*. sage, 1992.

[25] H. Deng, Q. Li, J. Huang, and C. Peng, "iCellSpeed: Increasing Cellular Data Speed with Device-Assisted Cell Selection," in *MobiCom*, 2020.

[26] H. Deng, K. Ling, J. Guo, and C. Peng, "Unveiling the Missed 4.5G Performance In the Wild," in *HotMobile*, March 2020.

[27] Q. Li and C. Peng, "Reconfiguring Cell Selection in 4G/5G Networks," in *ICNP*, 2021.

[28] 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.

[29] 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.

[30] G. Yang, X. Lin, Y. Li, H. Cui, M. Xu, D. Wu, H. Rydén, and S. B. Redhwan, "A Telecom Perspective on the Internet of Drones: From LTE-Advanced to 5G," *arXiv preprint:1803.11048*, 2018.

[31] 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.

[32] A. E. Garcia, M. Ozger, A. Baltaci, S. Hofmann, D. Gera, M. Nilson, C. Cavdar, and D. Schupke, "Direct Air to Ground Communications for Flying Vehicles: Measurement and Scaling Study for 5G," in *5GWF*. IEEE, 2019, pp. 310–315.

[33] J. Luo, P. Zhao, F.-C. Zheng, and L. Li, "Delay Evaluation for Cellular-Connected Drones: Experiments and Analysis," in *VTC2022-Fall*, 2022.

[34] 3GPP, "Work Item "Study on Enhanced Support for Aerial Vehicles"," June 2017, 3GPP TSG RAN Meeting #76, RP-171050.

[35] 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.

[36] 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.

[37] Q. Li, Z. Zhang, Y. Liu, Z. Tan, C. Peng, and S. Lu, "CA++: Enhancing Carrier Aggregation Beyond 5G," in *MobiCom*, 2023.

[38] Y. Liu, G. Guo, and C. Peng, "Demystifying Secondary Radio Access Failures in 5G," in *HotMobile*, 2024.