Measuring Fixed Wing UAS Networks at Long Range

Russell Shirey
rshirey@purdue.edu
Purdue University

Sanjay Rao
sanjay@purdue.edu
Purdue University

Shreyas Sundaram
sundara2@purdue.edu
Purdue University

ABSTRACT

In recent years, we have seen significant growth in Unmanned Aerial Systems (UAS), or drones, with many applications involving extensive communication from the UAS to the ground (e.g., sensor data transmission from the UAS). Despite being a different, and significantly more challenging environment than traditional networks, there exists limited understanding of UAS networking today, and the implications for applications. In this work, we conduct a long range measurement study of modern UAS over a 2-day period collected from a fixed wing UAS system. We consider both directional and omnidirectional antennas operating with Tactical Radios at long range distances. We pay particular attention to the edge cases, where performance is more challenging. Our results provide insights into how the dynamic nature of flight, plane orientation with fixed wing UAS, and the variability of wireless networks, impacts network performance.

CCS CONCEPTS

• **Networks → Network experimentation**; *Network range*; **Mobile networks**; Ad hoc networks.

KEYWORDS

UAS, UAV, Drones, Wireless, Networking, Throughput, Measurements, Tactical Radios

ACM Reference Format:
Russell Shirey, Sanjay Rao, and Shreyas Sundaram. 2020. Measuring Fixed Wing UAS Networks at Long Range. In *The 6th ACM Workshop on Micro Aerial Vehicle Networks, Systems, and Applications (DroNet'20), June 19, 2020, Toronto, ON, Canada.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3396864.3399697

1 INTRODUCTION

Recent technological advances have increased Unmanned Aerial Systems (UAS) usage at high rates and show no signs of stopping [6]. Many UAS applications involve sensor data collection and transmission of sensor data to an interested party on the ground, given that UAS can fly and go to places that humans cannot due to lack of physical access or as a safety precaution (e.g., areas impacted by disasters, law enforcement, etc.). The sensor data often has reliability requirements. Examples of such sensor data can include video with reliability requirements, Synthetic Aperture Radar (SAR), and Light Detection and Ranging (LiDAR). These sensors are often combined - putting significant demand on the network and requiring reliable networking to ensure important information is not missed due to dropped packets.

Since the UAS must travel to locations determined by the application requirements (e.g., providing mapping and oversight of an area to guide search and rescue missions), it is desirable to have long range options, and there is relatively limited freedom in placing UAS to optimize for connectivity. Today, civilian use of UAS in the United States (US) is typically restricted to limited distances that require visual line of sight [13] (typically, 1 Km). However, there is much interest in going to larger distances and regulations are beginning to support this worldwide [10].

Contributions: In this paper, we conduct a detailed measurement study from a fixed wing UAS over a 2-day period at distances exceeding traditional civilian regulations, through special approval. While there have been network studies with UAS, most have been conducted with multirotor UAS. Further, previous work focuses on experimentation at shorter distance with a mixture of 802.11 [1, 2, 8, 21] and LTE [5, 7, 11, 12, 17] networking focus. Our measurements focus on **long range distances** with **fixed wing UAS**. Additionally, our experiments are conducted in the context of **Tactical Radios** [3], which do not rely on pre-existing infrastructure. We are motivated to understand networking capabilities and limitations with UAS operating at long range distances, and with different omnidirectional and directional antenna configurations. Our measurements show how network throughput, and periods of dropout (where no data goes through), varies with flight path. Interestingly, the data indicates the **orientation of the plane relative to the ground node also significantly impacts network performance**. We pay close attention to situations on the edge of remote connectivity, where dropouts are common. Understanding the performance under these different scenarios will extend the usable distance of UAS, and greatly increase their applicability.

2 BACKGROUND

Fixed wing vs. multirotor UAS: There are two broad kinds of UAS - fixed wing and multirotor systems. Fixed wing systems are similar to traditional aircraft, with a central body and two wings. Multirotor systems are similar to a helicopter structure, with four (quad) or more rotors. We focus on fixed wing systems, since aerial coverage applications can benefit from faster and longer endurance UAS. However, fixed wing systems are much more difficult to characterize from a Radio Frequency (RF) standpoint, as they are not symmetrical in shape (leading to potentially varying wireless networking performance based on orientation), and typically must remain in motion in order to stay aloft.

Tactical Radios vs LTE/WiFi: Many aerial surveillance coverage applications cannot rely on pre-existing infrastructure, and hence require the entire wireless infrastructure on the UAS and the ground node. We refer to **Tactical Radios** [3] as radios that integrate the full infrastructure needed for communication into the

DroNet'20, June 19, 2020, Toronto, ON, Canada

ACM ISBN 978-1-4503-8010-2/20/06…$15.00
https://doi.org/10.1145/3396864.3399697

radio, allowing them to be interchangeable and operate in an ad-hoc manner. These radios utilize Mobile Ad-Hoc Networking (MANET) with multiple nodes, and can also operate in a point-to-point mode, with two dedicated radios for network traffic. This is relevant for UAS aerial coverage, with a radio on a UAS transmitting sensor data to another on the ground.

Tactical Radios have been used in the military for years and have since been adopted by Government and commercial entities, due to their capabilities. They are now commonly used to achieve goals in areas such as disaster relief [4], fighting wildfires [22], law enforcement [14], and crowd management and surveillance [15]. They have also been used to provide coverage of sporting events, such as the Super Bowl [19]. Due to their flexibility and applicability for use in long-distance UAS applications, we focus on Tactical Radios in this paper.

While there is recent interest in mounting LTE base stations on UAS [11], the technology is still under development. Furthermore, LTE has different infrastructure requirements based on whether the communication node is the master or slave. It is advantageous to have consistent networking equipment that is interchangeable and has both the master and slave functionality. This ensures that less backup equipment is needed, since it is not specialized, lowering the logistics footprint. Finally, the Tactical Radios we consider have much longer range than WiFi, presenting more options for aerial coverage at extended distances.

Types of antennas: UAS often operate with omnidirectional antennas, sometimes just referred to as **omni**, due to their constant movement. Omnidirectional antennas on the ground are the best for application flexibility, since they do not have to point at the UAS (and thus do not need a tracker). This enables the omnidirectional antenna to be mounted almost anywhere, even on a person, and moved around at ease. However, they have lower gain, resulting in less throughput than directional antennas. Directional antennas are larger and require a tracker to direct the antenna to the UAS, resulting in more equipment and setup time. There are trade-offs in either case and both are widely used in practice, depending on the application. Fig. 1 shows a comparison of the antennas we used.

3 MEASUREMENT METHODOLOGY

We seek to understand network performance in UAS settings, with the intention of enabling aerial surveillance over distances that stretch the limits of wireless connectivity. We characterize how network performance varies with distance, plane orientation, and antenna type. We present our measurement methodology, and discuss the data collected.

Regulatory Approval: UAS flight in the US is governed by the FAA [13]. Under civilian regulations, UAS are typically restricted to an altitude of 400 ft, and to distances that require visual line of sight for the duration of a flight (typically, 1 km). Given flight regulatory restrictions in the US, we collaborated with the Air Force Research Laboratory (AFRL) to accomplish relevant flight testing at distances exceeding FAA limits (recall we expect the FAA limits to increase over time (§1)).

Plane Selection: We flew a Martin UAV Bat-4, a representative fixed wing UAS in terms of size, weight, and speed, and appropriate for aerial surveillance activities. The Bat-4 flies 40-70 knots, depending on weight and wind conditions. The pilot controlled the airplane from the Ground Control Station (GCS) using a separate Command and Control (C2) link at a lower frequency that did not interfere with our data link.

Flight patterns: We introduce the flight terms **Distance, Slant Range, and Altitude**, and show them in Fig. 2. The plane orientation of **coming towards** (to the GCS) and **going away** (from the GCS) are also shown.

Figure 1: Antennas

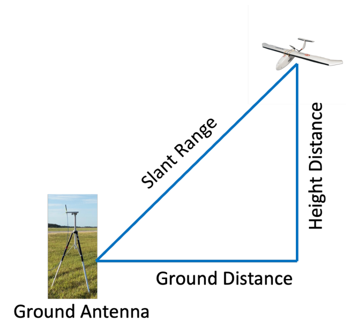

Figure 2: Orientation

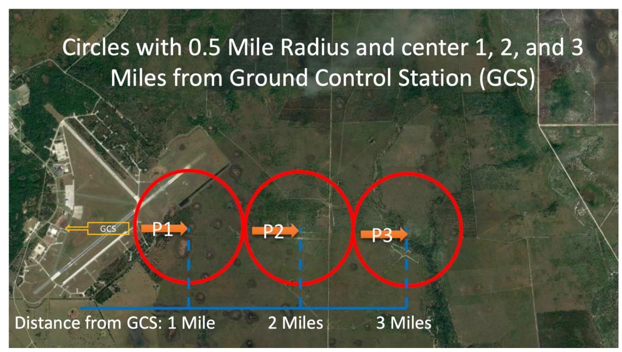

Figure 3: Flight plan and patterns

Circular orbits around a point are useful for recurring coverage of an area, and are common across many surveillance applications. For this reason, we mostly flew circle patterns to the east of the ground station, as shown and described in Fig. 3. We refer to the data collected from the circle orbits as **circ(k)**, with $k \in \{1, 2, 3, 4\}$, depending on the distance of the center of the circle from the GCS. We flew at 1500 ft cruising altitude Above Ground Level (AGL) and typically within 50-60 kts airspeed (taking about 160-180 seconds per circle orbit). With clear line of sight, the slant range determines the absolute distance between the radios. This distance affects the signal strength, and subsequent performance, of the network. Since slant range incorporates both ground distance and altitude into the calculation, we were able to test a wide variety of slant ranges while keeping altitude constant at a level that allowed for clear line of sight between the ground radio and the UAS.

Radio Selection: We used two Persistent Systems MPU4 Tactical Radios [18] tuned to S-Band frequency for testing. These radios provide seamless long-range layer 2 connectivity and support Internet Protocol (IP) traffic. We chose this configuration since the UAS sensor applications we consider require two radios in a point-to-point configuration: one on the UAS and the other on the ground.

Antenna Selection: We tested with an omnidirectional antenna on the plane and with both directional and omnidirectional antennas on the ground (see Fig. 1), since both are widely used in practice. We used a large directional antenna with 27 dBi gain on the ground (L-Com (HG2427)). This antenna aperture is 47.2 inches (in) x 35.43 in, offering significant gain but is very large and not practical for many applications. We used power ranging from 2W to 63mW, effectively lowering the gain to 12 dBi at 63mW, to be comparable to a smaller directional antenna, appropriate for more applications. We focus on this configuration throughout most of this paper, as the data collected with this configuration is more relevant to the edge of connectivity, given our flight distances. We use the term "directional antenna" to represent this configuration. A 12 dBi directional antenna is still larger than the omnidirectional and also requires a tracker to accurately point to the UAS. We used a small Haigh Farr 6130-4 omnidirectional blade antenna on the UAS.

Hardware setup: The UAS contains an air cooled payload bay on the bottom of the plane, where we stored a Raspberry Pi connected to a MPU4 radio to transmit and receive network data to the ground and store data on the UAS (using a 32 GB SD Card). On the ground, we used a laptop connected to a MPU4 radio to communicate with the UAS.

We scheduled flights on two different days to accomplish our goals. We flew all flight patterns for a given configuration in a single day to keep consistent results. We collected second-by-second throughput, latency, SNR, and location data with synchronized clocks. We collected TCP throughput information using iPerf [9]. We used TCP due to the dynamic and unreliable nature of the network operating at the extended distances we were flying, as well as the need for reliable sensor data in the applications we considered. We hosted the iPerf server on the Raspberry Pi in the UAS and ran the iPerf client on the ground laptop. We recorded latency via pings from the ground laptop to the Raspberry Pi. We recorded SNR and GPS information directly from the MPU4s.

4 DATA ANALYSIS

Our testing culminated in 6,245 seconds of throughput test time and 37 individual traces. The average throughput of each circular orbit test (0.5 Mile radius and center of circle 1 to 4 miles from GCS) is shown in Table 1.

Table 1: Circle Orbit TCP Average Throughput (Mbps)

Orbits With Center X Miles From GCS				
Config	**1**	**2**	**3**	**4**
Omni 2W	5.03	2.31	0.97	1.32
Dir 2W	15.40	13.20	12.00	10.80
Dir 500mW	14.60	13.00	11.70	10.10
Dir 125mW	14.60	11.50	9.08	6.92
Dir 63mW	10.80	6.90	4.92	3.26

As expected, we see that throughput decreases as distance increases (except for the anomaly of the omnidirectional 4 mile test, which will be explained later). Another observation is that the directional configurations all performed significantly better than the omnidirectional. This is expected because even the lowest gain directional configuration (63mW) is equivalent a 12 dBi antenna operating at 2W, much higher than omnidirectional. Table 1 can help guide applications for UAS operations with different configurations. In this paper, we focus on the 63mW directional and 2W omnidirectional configurations, as they present edge cases of network performance with challenges to consider when designing UAS applications.

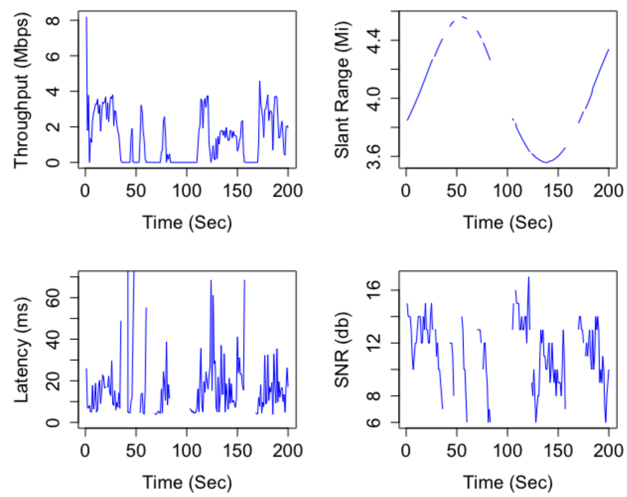

Figure 4: Omni circ(4) metrics

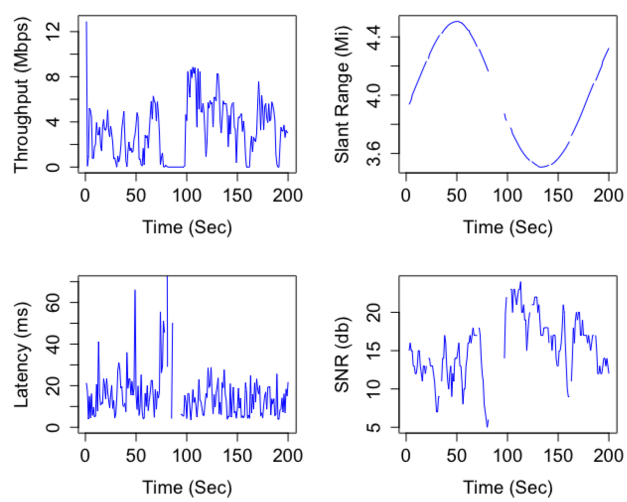

Figure 5: Directional circ(4) metrics

Network performance over flight path: Fig. 4 and Fig. 5 show a time series of the slant range of the plane, as well as the network metrics (throughput, latency, and SNR) for the circ(4) trace with omnidirectional and directional antenna configurations, respectively. The figures allows us to observe network performance, and how it varies with slant range. Dropouts are more prevalent at this distance, especially during the **coming towards** duration (from about 50 until 140 seconds for each circ(4) trace). Dropouts can be seen in Fig. 4 and Fig. 5 as sections where the throughput is 0 (top-left plot). Additionally, dropouts cause the lines to disappear in the other three plots. While the latency does not have a distinct pattern, **we observe several dropouts in all test cases.** Of note, we observe an increase in dropouts while the plane is in the **coming towards** orientation (especially between 80-100 seconds) of the traces; this qualitatively suggests that the orientation of the plane will have an impact on network performance. We will see that this impact is indeed present, and will quantify the differences in the subsequent subsections. The other omnidirectional traces and corresponding directional traces were qualitatively similar, with less dropouts as the range of distance decreased.

Fig. 6 and Fig. 7 show a time series of the same network metrics for the circ(1) trace with omnidirectional and directional antenna configurations, respectively. These tests have fewer dropouts than the circ(4) tests, due to being at a closer range. The SNR and throughput can clearly be seen as increasing as the slant range decreases, and decreasing as slant range increases.

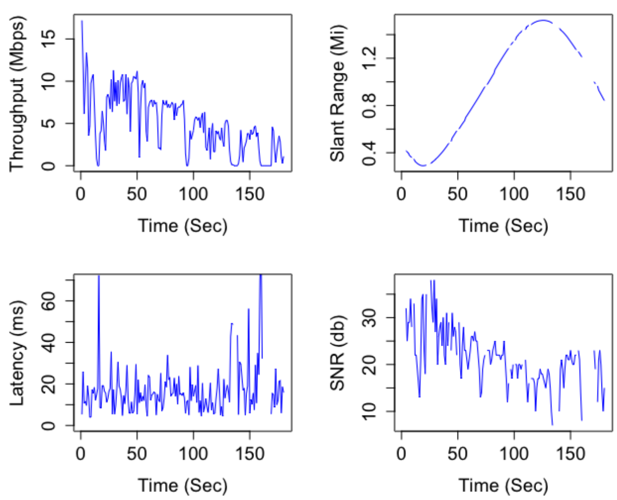

Figure 6: Omni circ(1) metrics

Antenna pattern and plane symmetry: The Bat-4 is not symmetrical and parts of the aircraft can interfere with wireless communication. As previously mentioned, this is common in fixed wing aircraft and optimal antenna placement often depends on the intended environment and flight patterns for the aircraft. Additionally, a closer look at the antenna pattern specification sheet for the UAS antenna shows weaker signal strength in the front of the antenna compared to the rear (by 1.5-2 dB). Both the UAS plane orientation and antenna pattern contribute to lower performance in the **coming towards** phase of the orbit, which is consistent with our data analysis. Next, we explore the impact of distance and orientation with regard to dropouts.

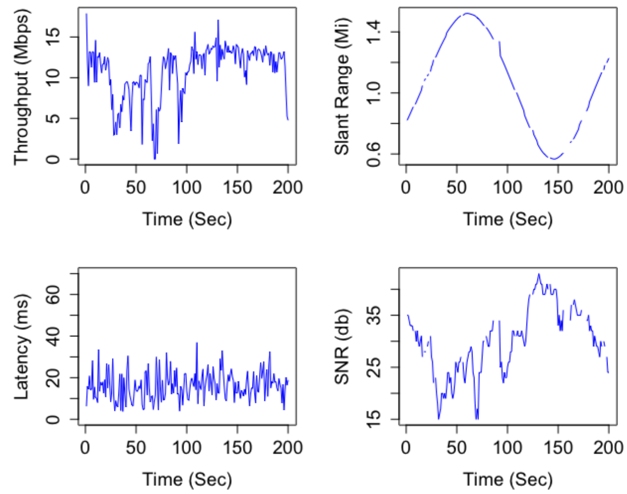

Figure 7: Directional circ(1) metrics

Dropout analysis: We define a **dropout period** as a period of at least one second in which the TCP throughput is zero. Dropout periods are typically caused by the SNR dropping below a certain threshold, which causes a temporary link failure in the radios. While dropouts are typical in wireless networks, especially at extended distances, we seek to further understand the dropouts based on distance and orientation.

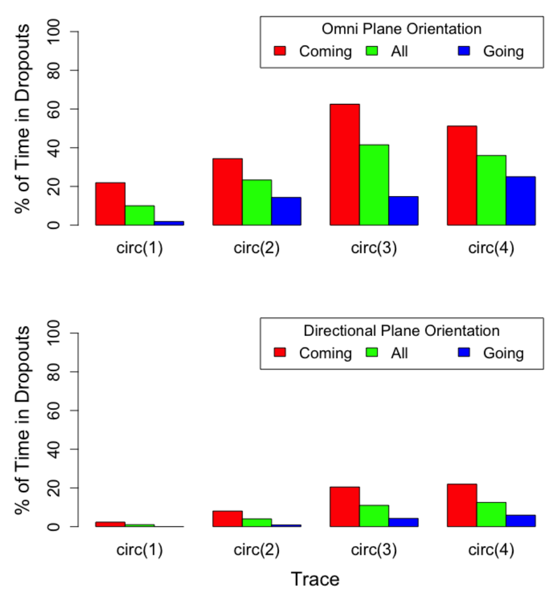

Figure 8: Percentage of Time in Dropouts

Fig. 8 shows the percentage of time in dropouts for each orbit, along with the orientation phases of the orbits. We observe

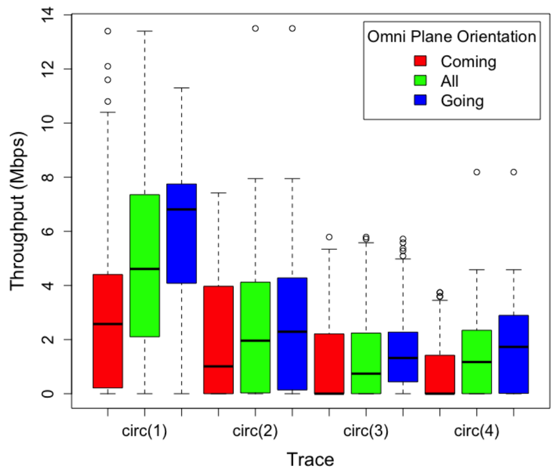

Figure 9: Omni throughput boxplots

the **coming towards** phase of the orbit has higher percentage of dropout time than the **going away** phase for both antenna types, validating the qualitative observations from the previous subsection. The omnidirectional tests experience significantly higher dropouts than directional, as expected. The percentage of time in dropouts increases sharply from circ(1) to circ(2) and again from circ(2) to circ(3), as distance increases. More interestingly, the omnidirectional test exhibits a decrease in percentage of time in dropout when moving from the circ(3) to circ(4) flight pattern. Upon review, we determined that this can be attributed to the time spent in each orientation during these tests. In particular, the UAS is **going away** for only 44% of circ(3), compared to 58% of circ(4). Each of these tests completed a full circle orbit and small part of another orbit for each (to ensure completeness, given slight variance in UAS speed). Each test started in a different position in the circle, resulting in different times spent in each orientation due to the extra portion of an additional orbit flown. We also considered average dropout duration, and found it generally increases with distance, and is higher for the **coming towards** direction.

Throughput analysis: We show boxplots of the throughput for each distance and orientation in Fig. 9. While there is a large variation in throughput across the traces, there are some clear trends that are consistent with previous findings of performance differences based on distance and orientation. We also analyzed SNR data, and it shows similar trends as throughput (SNR degrades with distance, and depends on plane orientation).

Distance: We notice in Fig. 9 that throughput generally decreases as distance increases, as expected. There is an exception to the trend as we move from circ(3) to circ(4), likely caused by the fact that circ(3) has additional time spent in the **coming towards** orientation, whereas circ(4) has additional time in the **going away** orientation, resulting in fewer dropouts.

Plane orientation: We also see in Fig. 9 that the third quartile in the **going away** phase is higher than in the **coming towards** phase across every trace. The first quartile is visibly higher for the **going away** phases in circ(1), circ(2), and circ(3), compared to the **coming towards** phases of those orbits. Due to the large number of dropouts, some of the quartile data is at or close to 0, making it difficult to directly compare.

Throughput time series analysis: We next investigate the correlation of throughput over time; this will indicate how much past bandwidth is a predictor of future bandwidth, with implications for different sensor data transmission algorithms. We use time series analysis testing to determine the correlation of the throughput data over time for each circle trace, with circ(1) and circ(4) omnidirectional and directional results shown in Fig. 10. The dashed lines represent the 95% confidence interval for an uncorrelated process; in other words, if the samples were uncorrelated over time, we would expect the sample autocorrelation at each lag to be inside the indicated bands with 95% confidence.

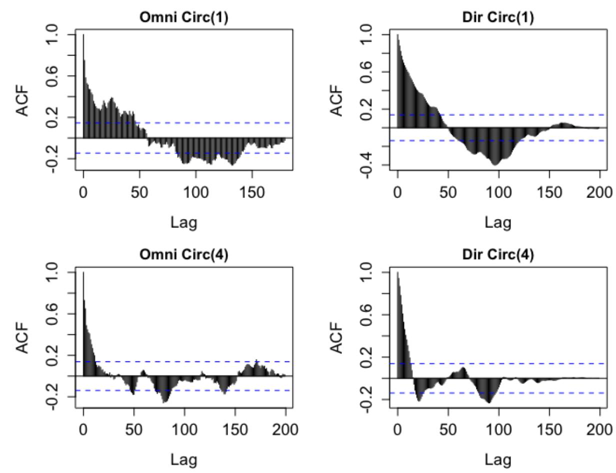

Figure 10: Autocorrelation of throughput

These tests show the throughput values are correlated over time. This is because the UAS moves in a seasonal pattern through the circle over time, resulting in performance differences based on the orientation and position of the plane. Interestingly, the autocorrelation plots also reveal the differences in the throughput due to the orientation of the plane. In particular, since the circular orbits take roughly 160-180 seconds to complete, we notice the anticorrelation of data occurring about halfway through, as a trend for all graphs. This represents the time-lag between the **coming towards** and **going away** orientations, agreeing with the results presented earlier. Additionally, the throughput is more closely correlated for longer lags at shorter distances and less correlated as the plane moves further away, due to increased throughput variability at further distances. Both trends agree with our earlier findings.

SNR throughput analysis: Fig. 11 shows a comparison of the throughput in relation to SNR for both antenna configurations. The results show a general increase in throughput as SNR increases, as expected. The throughput ranges for each SNR value are relatively similar for each antenna configuration, as expected. The overall throughput for directional is higher than omnidirectional because the flight is in a higher SNR range at the same distance, due to increased signal strength with a higher gain antenna.

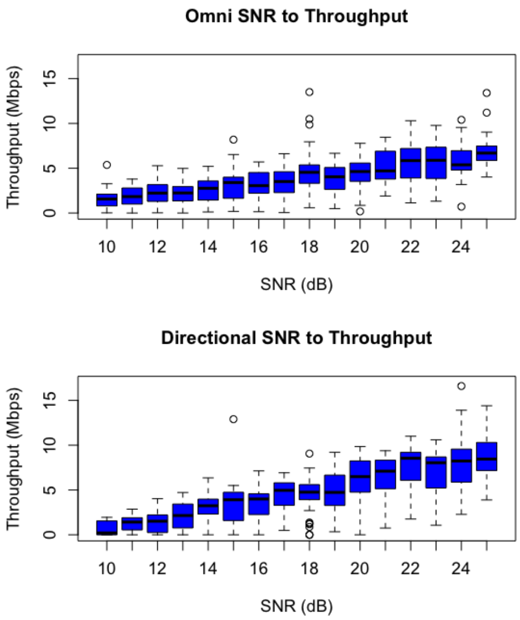

Figure 11: Omni and Directional SNR to throughput

5 RELATED WORK

There has been much recent interest in using UAS to extend Internet connectivity to remote locations [5, 11, 17]. The primary focus of these works is optimal positioning of the UAS to best serve the area [5], and how to best support LTE base station functionality on UAS [11]. There has been work with a focus on 802.11 fixed wing [1, 2] and multirotor [7, 8, 21] UAS networking, but at limited distances. Additionally, commercial In-Flight Communication (IFC) is characterized in [16], showing there is significant packet loss and throughput variation in such settings. A recent workshop paper [20] conducted a preliminary investigation of UAS for video streaming. Limited experiments with a multirotor and WiFi over short distances were provided. Finally, different types of Tactical Radios are tested with different transport protocols on the ground in [3]. In contrast to these works, our focus is on UAS networking for applications that require aerial surveillance coverage, often where it is not advisable to step foot close to the area being inspected. For these scenarios, it is important to enable acceptable performance over extended distances, and when the location cannot be optimized for connectivity. We focus on Tactical Radios rather than LTE settings, providing extended reach and no reliance on pre-existing infrastructure.

6 CONCLUSION

Our experiments explore **long range measurements** with **fixed wing UAS** using **Tactical Radios**. Table 1 shows the average throughput expected at different distance ranges with corresponding antenna configurations. However, as we explored, the network performance varies over the duration of the flight with **distance** and also with **plane orientation**, making networking with fixed wing UAS challenging. We focus on the edge scenarios, where dropouts are more common and performance is more unpredictable than at closer distances. Our results can be used to better tailor algorithms for delivery of critical aerial surveillance information, and improve the overall experience for end users. One implication is that we need to not only (i) consider the flight's distance, but also relative orientation; and (ii) we would need to explicitly consider dropouts in the application design. Our time series analysis shows potential benefits can be realized by taking advantage of the UAS flight pattern. In the worst case at the edge of connectivity, video delivery at certain resolutions may still be possible, but dropouts must be considered.

Acknowledgements: This work was partially supported by NSF Award ICE-T 1836889 and NSF CAREER award 1653648.

REFERENCES

[1] Mahdi Asadpour, Domenico Giustiniano, and Karin Anna Hummel. 2013. From ground to aerial communication: Dissecting WLAN 802.11n for the drones. In *WiNTECH 2013*. 25–32.
[2] Mahdi Asadpour, Domenico Giustiniano, Karin Anna Hummel, and Simon Heimlicher. 2013. Characterizing 802.11n aerial communication. In *MobiHoc 2013*. 7–12.
[3] Maggie Breedy, Peter Budulas, Alessandro Morelli, and Niranjan Suri. 2015. Transport protocols revisited. In *MILCOM 2015-2015 IEEE Military Communications Conference*. IEEE, 1354–1360.
[4] Ingrid Burrington. 2017. What Happens to the Internet After a Disaster. (2017). http://nymag.com/intelligencer/2017/10/what-happens-to-the-internet-after-a-disaster.html
[5] Ayon Chakraborty, Eugene Chai, Karthikeyan Sundaresan, Amir Khojastepour, and Sampath Rangarajan. 2018. SkyRAN: a self-organizing LTE RAN in the sky. In *ACM CoNext*. 280–292.
[6] Federal Aviation Administration (FAA). 2019. *FAA Aerospace Forecast Fiscal Years 2019-2039*. Technical Report.
[7] Samira Hayat, Christian Bettstetter, Aymen Fakhreddine, Raheeb Muzaffar, and Driton Emini. 2019. An experimental evaluation of LTE-A throughput for drones. In *DroNet 2019*. 3–8.
[8] Samira Hayat, Evşen Yanmaz, and Christian Bettstetter. 2015. Experimental analysis of multipoint-to-point UAV communications with IEEE 802.11n and 802.11ac. In *PIMRC 2015*. IEEE, 1991–1996.
[9] iPerf. 2019. (2019). https://iperf.fr
[10] Sue Mitchell. 2019. Unmanned aircraft flies first U.S. beyond-line-of-sight mission. (2019). https://news.uaf.edu/unmanned-aircraft-flies-first-u-s-beyond-line-of-sight-mission/
[11] Mehrdad Moradi, Karthikeyan Sundaresan, Eugene Chai, Sampath Rangarajan, and Z Morley Mao. 2018. SkyCore: Moving Core to the Edge for Untethered and Reliable UAV-based LTE Networks. In *ACM MOBICOM*. 35–49.
[12] Michael Nekrasov, Ryan Allen, and Elizabeth Belding. 2019. Performance Analysis of Aerial Data Collection from Outdoor IoT Sensor Networks using 2.4 GHz 802.15.4. In *DroNet 2019*. 33–38.
[13] Code of Federal Regulations (CFR). 2019. 14 CFR: Aeronautics and Space. (2019).
[14] Department of Homeland Security (DHS). 2017. *MANET Operational Field Test with New York City Police Department Emergency Service Unit*. Technical Report.
[15] Department of Transportation (DoT). 2018. *FHWA White Paper On Mobile Ad Hoc Networks*. Technical Report.
[16] John P Rula, James Newman, Fabián E Bustamante, Arash Molavi Kakhki, and David Choffnes. 2018. Mile high wifi: A first look at in-flight internet connectivity. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 1449–1458.
[17] Karthikeyan Sundaresan, Eugene Chai, Ayon Chakraborty, and Sampath Rangarajan. 2018. SkyLiTE: End-to-End Design of Low-Altitude UAV Networks for Providing LTE Connectivity. (2018).
[18] Persistent Systems. 2016. MPU Specification Sheet. (2016). http://www.persistentsystems.com/pdf/MPU4_SpecSheet.pdf
[19] Persistent Systems. 2017. Super Bowl Broadcasters Leverage MPU5 For Live Shots. (2017). https://www.persistentsystems.com/super-bowl-broadcasters-leverage-mpu5-for-live-shots/
[20] Xiaoli Wang, Aakanksha Chowdhery, and Mung Chiang. 2016. SkyEyes: adaptive video streaming from UAVs. In *HotWireless 2016*. ACM, 2–6.
[21] Evşen Yanmaz, Robert Kuschnig, and Christian Bettstetter. 2013. Achieving air-ground communications in 802.11 networks with three-dimensional aerial mobility. In *INFOCOM 2013*. IEEE, 120–124.
[22] Mark Zaller. 2017. Fire View Infrared Downlink (FirstNet). Wildland Fire Common Operating Picture and Interactive Drawing. (2017).