

Carnegie Mellon University

From the Selected Works of Bob Iannucci

June, 2017

Holes-in-the-Sky: A Field Study on Cellular-Connected UAS

Ervin Teng, *Carnegie Mellon University*

João Diogo Falcão, *Carnegie Mellon University*

Robert A Iannucci, *Carnegie Mellon University*



Available at: <https://works.bepress.com/bob/30/>

Holes-in-the-Sky: A Field Study on Cellular-Connected UAS

Ervin Teng, João Diogo Falcão, and Bob Iannucci
Department of Electrical and Computer Engineering
Carnegie Mellon University

NASA Ames Research Park, Building 23 (MS 23-11), Moffett Field, CA 94035
{ervin.teng, joao.diogo.de.menezes.falcao, bob}@sv.cmu.edu

Abstract—Small unmanned aerial systems (UAS) require constant, safety-critical connectivity for telemetry, command-and-control, and collision avoidance. Today, dedicated, short-range pilot-to-aircraft links provide this connectivity for UAS operation. For UAS operating in fleets and beyond line-of-sight, a robust multi-transmitter network to provide connectivity over a wide area will be needed. However, networks that could serve this purpose, such as the ubiquitous broadband cellular networks, were planned and deployed for terminals on the ground. Hardening multi-transmitter networks for aerial use remains an open problem.

In this paper, we demonstrate through field measurement that a typical cellular deployment could result in low-coverage areas for UAS—what we call the “hole-in-the-sky” phenomenon. Furthermore, many of the propagation models and assumptions commonly used in terrestrial network planning fail to accurately predict aerial signal strength. From first principles, we identify and model the predominant contributors to path loss, and form a combined propagation model that more accurately reflects reality for the tested scenarios. Motivated by this study, we identify a new research direction towards avoiding holes-in-the-sky during flight.

I. INTRODUCTION

Small unmanned aircraft systems (UAS) have been steadily dropping in cost, due to the proliferation of low-cost microelectromechanical systems-based (MEMS) sensors, global positioning system (GPS) receivers, and mobile processors. They are finding use in a number of commercial applications, such as aerial photography and precision agriculture. We can imagine a future where such systems are used ubiquitously for a variety of aerial applications, such as goods delivery or wide-area surveillance and monitoring.

Being unmanned, UAS require constant connectivity. A wireless link to the ground is crucial for telemetry information, live video or sensor feeds, sending navigational commands, and manual override. In addition, these wireless links can be used to share position and trajectory information for collision avoidance. These links are safety-critical, as losing them can mean loss of control and a possible crash. Today, civilian and commercial UAS rely on a single point-to-point wireless link between the aircraft and an operator on the ground for this functionality. But fleets operating beyond the line-of-sight of any single operator will require a multi-transmitter network that is relatively low cost, high-bandwidth, cloud-connected, and reliable.

Cellular networks, particularly with the Long-Term Evolution (LTE) standard, have the requisite wide area coverage, low cost of integration, packet-switched data capability, coding schemes and network infrastructure designed for moving clients, and sub-100ms latency [7]. Current work on re-evaluating cellular network architecture [8] would naturally benefit UAS clients, with their high dependability requirements and unique bandwidth needs. Hardening and adapting today’s cellular infrastructure for UAS offers a more practical alternative to a purpose-built secondary network.

However, current cellular networks are first and foremost designed for revenue-generating customers on the ground, and the above-ground coverage is not well characterized. The radio frequency (RF) propagation models commonly used to plan and evaluate such networks were built on datasets collected on the ground [6][3], and often do not allow the receiver to be higher than the transmitter in their formulation. Networks deployed for terrestrial use could produce unexpected coverage “holes-in-the-sky,” which could span tens to hundreds of meters and present a hazard for connectivity-reliant drones. In order to mitigate this hazard, drones will need an awareness of the channel state at their location, and the ability avoid them while navigating.

In this paper, we seek to evaluate the possibility of using cellular and other wide-area networks for command and control of UAS, through a series of field studies, and the feasibility of using simple modeling for real-time prediction of coverage.

This paper’s contributions are as follows:

- Demonstration and modeling of the hole-in-the-sky phenomenon, using field measurements. We depict potential cellular deployment scenarios and study the resulting coverage holes as seen by a UAS.
- An evaluation of a compositional model consisting of two-ray ground reflection, incursion of obstacles into the Fresnel zones, and antenna pattern for common scenarios faced by cellular-connected UAS.

Furthermore, these results suggest that real-time estimation of path loss by a UAS in flight, based on simple geometrical knowledge of the scenarios, could be an effective strategy to avoiding holes-in-the-sky.

This paper is structured as follows. Section II provides an overview of propagation modeling techniques. Section III de-

scribes the design and characterization of the RF measurement system. Section IV describes the experimental setup, Section V develops a model based on the collected data, and Section VI applies these models to several test scenarios. Section VII shows where holes-in-the-sky can form, and how we might mitigate such effects in the future. Section VIII presents related studies, and we conclude with Section IX.

II. RF PROPAGATION MODELING

This section presents an overview of path loss modeling, and the mathematical basis for modeling two-ray ground reflection and diffraction around obstacles.

A. Signal Strength and Path Loss

In planning wireless networks, we must ensure that the receiver receives enough RF power to decode the signal. The *received power* P_r , measured in dBm, seen by a receiver, is given by Equation 1. P_r is a function of the antenna gains G_t and G_r , transmit power P_t , system losses L_{sys} and path loss L_{prop} [10].

$$P_r = P_t + G_t + G_r - L_{sys} - L_{prop} \quad (1)$$

P_t , G_t , G_r , and L_{sys} are factors of the particular radio system in use; L_{prop} varies with environment. For instance, in the idealized case without obstacles or terrain, L_{prop} can be approximated by L_{fs} , as given by Equation 2 [10].

$$L_{fs}(\text{dB}) = 32.4 + 20\log(d(\text{km})) + 20\log(f(\text{MHz})) \quad (2)$$

In cases where the receiver is close to the transmitter and within line-of-sight, free-space path loss may be a good approximation of propagation and can be used to predict signal strength. In most cases, however, other factors contribute to the path loss. The following sections present ways to account for several of these effects.

B. Path Loss Models

Models for path loss fall into two main categories. *Empirical* models, such as the Okumura-Hata model [6] for cellular and land mobile applications, are created from experimental data. While general and only accurate on a macro scale, their ease of use and computationally light nature make them popular. *Deterministic* models are based on environmental information and an approximation of physics. They grow in complexity with the environment—for instance, performing RF ray-tracing on a complex urban scene is computationally very expensive. In between are models that attempt to balance usability and accuracy, such as the Walfisch-Ikegami model [5], which takes into account not only the properties of the receiver and transmitter but also average building heights and street widths.

In [15], we examined the Walfisch-Ikegami model for UAS use; the model's operating conditions did not apply to a UAS client being served by a cellular tower. The same is true of other empirical and semi-empirical models [6], as they only allow for receivers a few meters off the ground. In this work, we return to first principles and basic deterministic models.

C. Two-ray or Plane-earth Propagation

In real-world propagation, we must account for not only the loss on the single path in Equation 2 but also the contribution of other paths, *i.e.* *multipath*. In simple, unobstructed scenarios, the main reflected path comes from the ground and sums with the direct path, creating interference from the two paths' phase alignment that is distance and frequency dependent. Figure 1 shows this geometry.




Fig. 1: Geometry of two-ray reflection.

Experimental results by Sommer *et al.* [14] show a good correlation between field measurements in an open environment and this *two-ray interference model*. We rewrite Equation 2 as:

$$L_{fs} = 20 \log\left(\frac{4\pi}{c} df\right) \quad (3)$$

where f represents the carrier frequency of the transmitted wave.

In order to account for the reflected path, we introduce a reflection term

$$\Gamma_{\perp} = \frac{\sin \theta_i - \sqrt{\epsilon_r - \cos^2 \theta_i}}{\sin \theta_i + \sqrt{\epsilon_r - \cos^2 \theta_i}} \quad (4)$$

And a phase difference

$$\gamma = 2\pi \frac{d_{ref} - d_{los}}{\lambda} \quad (5)$$

If we assume that the ground is flat between the transmitter and receiver, we can compute the reflected path length as

$$d_{ref} = \sqrt{d_{ground}^2 + (h_{rx} + h_{tx})^2} \quad (6)$$

where h_{rx} and h_{tx} are the heights of the receiver and transmitter.

We assume that the RF energy is predominantly vertically polarized, and do not consider the reflection terms for non-vertical polarization. We make this approximation as both our transmit antenna and receive antenna are vertically polarized.

We can amend our free-space path loss equation to include the reflection effects in Equation 7.

$$L_{tr} = 20 \log\left(\frac{4\pi}{\lambda}\right) - 20 \log\left(\left|\frac{1}{d_{los}} + \frac{\Gamma_{\perp} e^{i\gamma}}{d_{ref}}\right|^{-1}\right) \quad (7)$$

Figure 2 shows a plot of Equation 7 with the transmitter 14m and the receiver 2m above ground level (AGL).

We denote a turnover distance d_t by Equation 8.

$$d_t = \frac{4\pi}{\lambda} h_{tx} h_{rx} \quad (8)$$




Fig. 2: Two-ray interference path loss, showing turnover distance d_t .

Beyond this turnover distance, the path loss increases with the fourth power of distance and does not fluctuate. Because d_t , with h_{rx} at 1-2m for a terrestrial receiver, is small compared to the radius of a cell site, many propagation models use Equation 9 as the base path loss in absence of obstacles. This approximation is often referred to as the plane-earth or “two-ray model” of path loss.

$$L_{PE} = 40 \log_{10}(d_{ground}) - 10 \log_{10}(h_t^2 h_r^2) \quad (9)$$

However, in the case of UAS links, where h_{rx} can be 40m or more, d_t is increased by many times, and we must reconsider the use of this approximation to describe two-ray propagation.

D. Diffraction Loss from Obstacles

Like all electromagnetic waves, RF waves are not completely blocked by objects in the line-of-sight path. Rather, obstacles cause blockage of the Fresnel zones, *i.e.* successive regions of constructive and destructive interference around the line-of-sight path. Energy from unobstructed Fresnel zones still propagates to the other side [11].

To compute the expected electric field strength of an electromagnetic wave after propagating through an aperture in an infinitely-thin screen, we can use the Fresnel-Kirchoff equation [10][11]. We examine the case where the aperture is open in all directions except the bottom, *i.e.* knife-edge diffraction. For each obstacle, we can compute the diffraction parameter ν_i , given by the equation

$$\nu_i = c_i \sqrt{\frac{2(d_1 + d_2)}{\lambda d_1 d_2}} \quad (10)$$

where d_1 and d_2 are the distances from the Tx and the Rx to the obstacle, respectively, and c_i is the negative distance from the Rx-Tx path to the top of the obstacle (Figure 3).




Fig. 3: Geometry of obstruction and RF path.

We can then model the normalized electric field, *i.e.* the ratio of the field strength after the aperture to the one before, using the Fresnel integral [11]:

$$\frac{E_d}{E_o} = F(\nu_i) = \frac{1+j}{2} \int_{\nu_i}^{\infty} e^{-\frac{j\pi t^2}{2}} dt \quad (11)$$

And subsequently, the resulting path loss as:

$$L_{diff} = -20 \log |F(\nu_i)| \quad (12)$$

Figure 4 is a plot of L_{diff} versus ν_i . Note that the value of this function for $\nu_i < -0.8$ (green line) is close to 0, a common first-order approximation [4].




Fig. 4: Path loss vs. ν_i . Function is well-known and independent of environmental parameters.

We can then define our *total* path loss after the knife-edge to be:

$$L_{prop} = L_{diff} + L_{base} \quad (13)$$

where L_{base} is the path loss in the absence of the obstacle.

E. Extension to Wideband Signals

In the prior equations, we made the assumption that the signal consists of a single carrier frequency, while signals, such as LTE, that use orthogonal frequency-division multiplexing (OFDM) will have multiple carriers at different frequencies. However, the difference in behavior of these carriers is typically small for simple scenarios. For example, in free-space, for two carriers 5 MHz apart at 900 and 905 MHz, the loss would vary by less than 0.05 dB. Similarly, in the two-ray scenario shown in Figure 2, the depth of the deepest null varies by only 0.3 dB over the same 5 MHz frequency range. In this paper, we make the modeling approximation that the signal consists of a single frequency. In more complex multipath environments where frequency-selective fading is very narrow, multi-carrier schemes will be *more* robust [10] than a narrowband signal, and would suffer from performance degradation rather than loss of signal.

F. Computing Path Loss from Signal Strength

In studying the path loss caused by factors in the environment, we must first remove the effects of the measurement system. We can do this by taking two measurements of received RF power, P_{d_0} and P_{d_1} , at two distances from the transmitter in the same direction. Substituting Equation 1, we see that:

$$P_{d_1} - P_{d_0} = L_{d_1} - L_{d_0} = L_{d_0 d_1} \quad (14)$$

$L_{d_0 d_1}$, the path loss along the path between distance d_0 to distance d_1 , is independent of antenna gain, transmit power, and system losses, and will be the subject of our path loss studies. For consistency, we fix d_0 . The following section will describe the instrument used to measure P_{d_0} and P_{d_1} , and the instrument’s characterization.

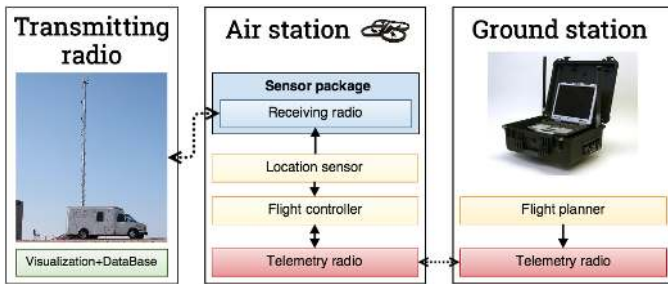


Fig. 5: System architecture [15].

III. INSTRUMENT DESIGN AND CHARACTERIZATION

We designed our aerial probe for received signal strength around the following principles:

- Measurements at any given point in space should be independent of the aircraft's attitude.
- Signal strength readings should be referenced to a lab spectrum analyzer to within 3 dB.
- *Path loss* should be measured within 5 dB, and independent of the antenna gain and transmit power of the instrument.

A. System Design

Our measurement system consists of a transmitter-under-test, an aerial probe, and a ground station which records data sent from the aerial probe. A diagram of the system is shown in Figure 5.

Aircraft and Flight Control. In order to carry our sensor payload, we selected a quadcopter for its maneuverability and ability to hover. The sensor package is mounted to the base of the craft (Figure 6a). An airframe with retractable landing gear allows the conductive carbon fiber gear to be moved away from the antenna, minimizing its effect on the antenna's pattern. The flight controller is based on the ArduCopter¹ software stack, and supports autonomous GPS-based navigation as well as assisted manual flight. An SBAS²-capable GPS is used for providing navigational as well as sensor localization. Figure 6b shows the aircraft in flight.

Sensor Package. For size and weight reasons, we use a pair of small narrowband radios operating with a center frequency of 909 MHz. based on the HopeRF chipset. These radios operate on a simple time-division multiple access (TDMA) scheme and report an 8-bit received signal strength indicator (RSSI) value, for both radios at both ends, every second. It is important to note that the HopeRF radio has a bandwidth of 105 kHz, while cellular signals tend to have a much wider bandwidth. We maintain the single-frequency assumption as stated in Section II-E. A vertical 1/2 wavelength antenna is mounted on a servo-driven gimbal to keep the antenna vertical as the aircraft tilts.

Ground Station. RSSI data and location information are sent over the sensor radio link to the ground station, where

¹ ArduCopter. <http://copter.ardupilot.com>

² satellite-based augmentation system

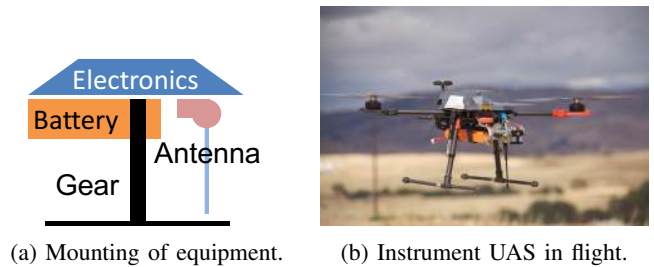


Fig. 6: Aerial instrument.

it is stored in a database. A web interface allows viewing of collected data in real-time.

Transmitting Radio. The transmitter-under-test was housed in a mobile communications van (Figure 5) equipped with a pneumatic mast. Atop the van's mast was either a vertical omnidirectional antenna (L-com HGV-906U, 30° vertical beamwidth, 6 dBi gain), or a directional sector antenna (L-com HG913P, 20° horizontal 15° vertical beamwidth, 13 dBi gain, Figure 7). With the mast extended, the antenna stood 13.9 m above ground level. We sent test packets from the transmitter-under-test at 20 dBm (measured at the transmitter's output).

Transmitter Properties	Value
Transmit Power	20 dBm
Antenna Height (AGL)	13.9 m
Transmit Frequency	909 MHz
Max. Antenna gain (HGV-906U)	6 dBi
Max. Antenna gain (HG913P)	13 dBi



Fig. 7: Sector antennas.

B. Instrument Characterization

We characterized the positional error of the instrument by measuring the GPS drift and accounting for the refresh rate of both the radio and the GPS. Subsequently, we mapped RSSI readings to dBm by feeding an attenuated signal into the radios, taking RSSI readings, and then measuring actual received power from a spectrum analyzer. In-flight, we also measured the effect of yaw angle on received power.

In addition, we conducted two system-level tests in-flight. In the first, we determined the in-flight effects of aircraft velocity on measurement. In the second, we determined the antenna gain pattern of the coupled drone-transmitter system to use in the computation of (antenna independent) path loss.

1) *Measurement Consistency with Velocity:* In this section, we examine the instrument's repeatability and invariance to velocity and direction. We assume that for a fixed transmitting antenna, fixed transmitter power, and fixed environment, the signal strength at any particular point in space remains fixed,

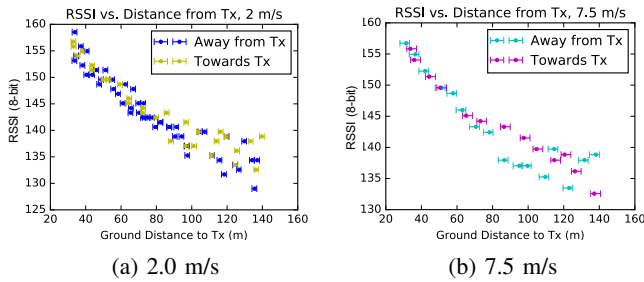


Fig. 8: Linear runs at different velocities.

and that variability in measured signal strength is caused by the measurement device itself. We flew the UAS horizontally in a straight line from the transmitter to a distance of 400m and back, with varying velocities. Figure 8 shows the recorded values from this test at velocities of 2.0 and 7.5 m/s.

We observed little velocity-dependent variation. However, at lower velocities, we noted a *direction*-dependent variation, at the same magnitude (2 dB) as the front-to-back ratio observed during in-air rotational tests. This difference diminishes as velocity increases and the angle of attack moves the battery out of the antenna path. To mitigate the effects of the yaw angle on our readings, we flew all measurements either towards or away from the transmitter, and maintained velocity at or above 10 m/s.

2) *Coupled System Antenna Pattern*: Using the omnidirectional antenna, we can factor out the vertical elevation pattern of both the transmitting antenna and the receiving antenna to study path loss independently. Referring to Equation 14, one way to do this is to take a reference measurement close to the transmitter for each radial direction from the transmitter pointing towards a data point. In order to compute radial path loss from the antenna at a constant d_0 , we need measurements in a sphere around the antenna, beyond the near-field ($d > \frac{2D^2}{\lambda}$ [11], where $D = 0.6$ m).

However, for practical reasons we could only fly a cylindrical pattern around the antenna (Figure 9a). We flew the craft around the transmitting antenna at $d_0 = 14.8$ m, and at 15 different altitudes between 10 m and 60 m AGL. We then recomputed the received power as if it was on the sphere. Rephrasing Equation 2, we can say that the received power in dBm is given by $P_r = C - 20\log_{10}(d)$, where C is a constant independent of distance. We can compute C given the current radial distance to the cylinder $d = d_{circ}$, and subsequently solve the same equation for the received power as it would be if $d = d_0$.

We then fitted a 3rd-order polynomial to the relationship between elevation angle and received power at d_0 . These signal strengths can then be treated as the received power at d_0 in Equation 14. Figure 9b shows the computed received power vs. elevation angle, and the fitted curve used in future computations of path loss. The strong lobe to the side, and deep nulls above and below, are consistent with expectation for a monopole ground plane antenna.

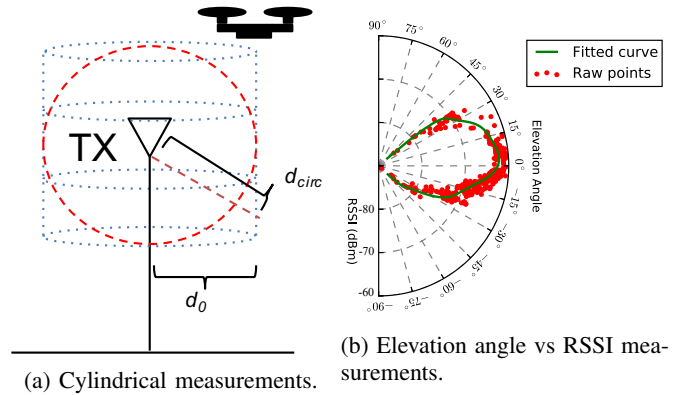


Fig. 9: Determining antenna gain pattern. The UAS was flown in a cylinder at a far-field radius to establish antenna-gain-independent reference power levels for use in subsequent path loss measurements.

C. Summary of Characterization

Through lab and in-flight characterization, this system was determined to have the following properties, at a flight velocity of 10 m/s. Error values are presented as the 95% prediction interval of the antenna pattern and RSSI-to-dBm functions.

Property	Value ($v = 10$ m/s)
Measurement Accuracy	± 2.15 dB
Path Loss Accuracy	± 3.90 dB
Positional Error	± 4.36 m

One limitation of this instrument is that the positional error is greater than a wavelength ($\lambda = 0.33$ m). But the likely cause of coverage *holes*, large-scale fading, happens at geometries greater than this position error.

IV. FINDING HOLES-IN-THE-SKY

Using our instrument and transmitter, we performed experiments at a Joint Inter-agency Field Experimentation (JIFX) event, held by the Naval Postgraduate School at Camp Roberts, CA. We conducted propagation studies in two scenarios: An open field, and a mock village (Figure 11) with buildings and other obstacles.

At the open field, we flew a ladder pattern—linear paths at various altitudes, out to 400 m from the transmitter. Figure 10a shows the flight path.

Figure 12 is a panoramic view from the antenna toward the primary obstacles in the mock village. Of particular note are the marked buildings, which we will refer to as “Hotel,” “Church,” and “House.” In order to gather information about the effect of obstacles on RF paths in this environment, we flew the craft above the buildings in a radial pattern (Figure 10) out from the location of the transmitter to 400 m ground distance, at target altitudes of 40, 60, 80, and 100 m above ground. We also gathered radial data behind and around buildings.

V. MODELING HOLES-IN-THE-SKY

Using the data from Section IV, we show that the two-ray interference model is a better fit than the purely logarithmic



(a) Linear ladder pattern. (b) Radial ladder patterns.

Fig. 10: Experimental flight patterns for open field and mock village, respectively. Black circle marks location of the transmitter.



Fig. 11: Mock village at Camp Roberts, CA.

models identified in [15], that the diffraction around obstacles are predictable to a first order, and that the sector antennas typically used in cellular deployments could lead to aerial coverage problems.

A. Open Field Propagation

We use data collected at the airfield to create a baseline for path loss. Figure 13 shows path loss from a single radial flight away from the transmitter at 35m above ground level and compares the measured data with the predicted curve from free-space path loss. On average, we observe the roll-off predicted by free-space path loss. The two dotted lines above and below the data points represent the 95% prediction interval of the polynomial fit used to map elevation angle to reference RSSI.

We observe that while, on average, Equation 2 alone could be used to approximate this path loss, there is a second strong effect—the interference from a second ray reflected by the ground. The relative phase of these two rays causes the signal to fluctuate with distance from the transmitter. We can model this additional effect with the formulation in Section II-C.

Adjusting for the experiment-specific parameters, we use Equation 7 to compute path loss for the same distance ranges we measured. We initially begin with a σ_r of 4, which from literature [11] would have been a *low estimate* for the relative permittivity for dry, flat ground. However, using this parameter to generate our simulated results, we found the phase-interference oscillation of signal strength to be much greater than in our experimental data. We then perform a least-squares fit of the simulated two-ray interference path loss with the experimental data. In this fashion, we obtain a σ_r value of 1.037. This reflection coefficient matches the

experimental results from [2] and [14]. Figure 14 shows the resultant computed path loss using the fitted two-ray model along with the measured values. The resultant standard error of the fitted model is $\sigma_{est} = 1.43$ dB. We will use this fitted model as the base path loss in subsequent analysis.

Coming back to the issue of turnover distance, we observe that, at a flight altitude h_r of 35 m and a transmitter height of 13.9 m, Equation 8 predicts a turnover distance of 5.9 km. Higher altitudes would only increase this distance. Consequently, the plane-earth path loss model (Equation 9) is a poor fit for aerial links. For UAS, we must include two-ray interference in our model.

B. Obstacles and Diffraction Loss

We now consider how obstacles in the RF path affect path loss, using data from within the mock village environment from Figure 12. Measurements taken behind the Church (Figure 16) do not correlate well with prediction using only free-space propagation and two-ray effects. When flying below the roof height, we expect a significant loss from the building’s incursion into the Fresnel zones. As the building has a slim profile relative to the flight distances, we approximate the building as a knife-edge and apply the knife edge diffraction method described in Section II-D. Figure 15 shows our simplified representation of the church structure (as a knife-edge) and roll-off of the terrain behind the church (as a single step change in ground height).

Equation 13 is used to compute the total path loss L_{prop} after the obstacle. As L_{base} , we use the two-ray model with the relative permittivity computed in Section V-A. Equations 10, 11, and 12 are used, with the appropriate geometry, to compute L_{diff} .

Figure 16 compares the measured data with L_{base} and the total loss $L_{prop} = L_{base} + L_{diff}$. Modeling only L_{base} , the two-ray path loss, results in a $\sigma_{est} = 7.46$ dB and $E_{max} = 19.3$ dB. Taking into account L_{diff} results in a $\sigma_{est} = 3.42$ dB, and a maximum absolute error of $E_{max} = 8.52$ dB, a marked improvement. In an attempt to better model the knife-edge diffraction scenario, we apply the *four-ray model*, where rays to not only the Tx and the Rx but to their mirror images are accounted for. We see a marginal improvement to $\sigma_{est} = 2.21$ dB—but a large increase in computational complexity.

Given the Rx/Tx heights, Rx/Tx distance, and the location and height of the obstacle, we are able to model our measured path loss for the single-obstacle case to 10 dB by combining the two-ray model and Fresnel-Kirchoff diffraction.

C. Coverage from a Directional Antenna

Prior experiments described here and in [15] have used a vertical, omnidirectional antenna, so that flight paths radially out from the transmitter are comparable in analysis and the effects of the environment can be better studied. Typical commercial cellular deployments, however, use highly directional sector antennas, which have a wide horizontal beam width (120° or more), a narrow vertical beam width (20° or less), and a high gain of 10 dBi or more. These antennas direct RF power




Fig. 12: Panoramic view of the environment from the transmitting antenna.

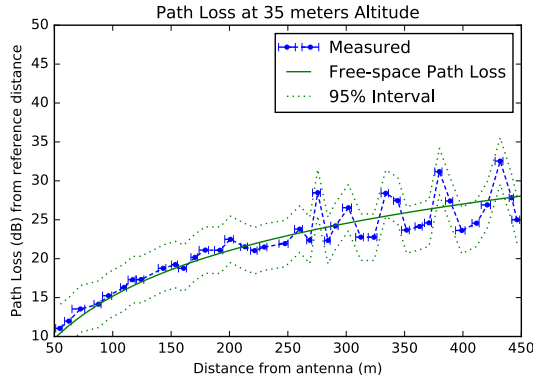


Fig. 13: Path loss data from a single flight at 40 m AGL. Interference pattern clearly visible. Dotted lines show 95% confidence interval of path loss curve (Section III-B2).

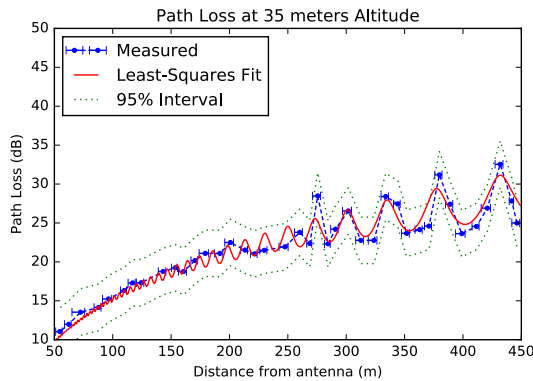


Fig. 14: Computed effect of ground ray interference with measured values. Standard error $\sigma_{est} = 1.43$ dB.

where needed, and allow for cell splitting—multiple sectors on one tower. To focus coverage towards the ground, these antennas are tilted slightly downwards—detrimental to aerial coverage. How would these antennas work for UAS?

Using the sector antenna in place of the omnidirectional one, we conducted measurements at the open airfield in a similar fashion to Section V-A with the sector antenna pointed down the runway, flying an aerial ladder pattern at various altitudes. We expect the strongest signal strength in front of the antenna. Figure 17 shows the measurements for this experiment, and Figure 18a shows signal strength along the highest altitude path. We note in red significant drops in signal strength at various vertical elevation angles from the transmitting antenna.




Fig. 15: Approximation of Church and terrain change as a knife-edge plus a single, discontinuous change in ground height.

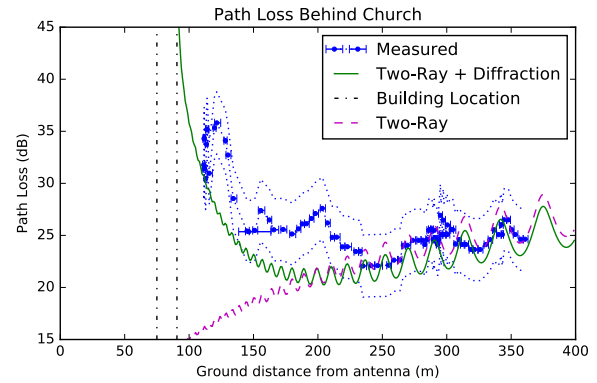


Fig. 16: Comparison of simulated path loss from the knife-edge and two-ray ground reflection against data collected behind Church. Prediction has σ_{est} of 3.42 dB.

This corresponds to the published radiation pattern of this antenna (Figure 18b), which shows four deep nulls in the vertical radiation pattern. Measured signal strength at the trough of a null is 14 dB below the signal measured just outside of a null. For receivers above the antenna’s forward main lobe, there will be significant drops in signal strength, especially directly above the transmitter. These antenna nulls create additional coverage holes that UAS can unexpectedly enter during flight. In the worst case, a UAS traveling in parallel to a null could remain in it for the entire time spent in that cell.

To summarize, we identified three main effects, beyond free-space path loss, in our measurements. We have modeled two-ray interference and diffraction loss around an obstacle to within 10 dB of measured values, and demonstrated that the antenna pattern of a sector antenna can produce additional, albeit predictable, coverage holes.

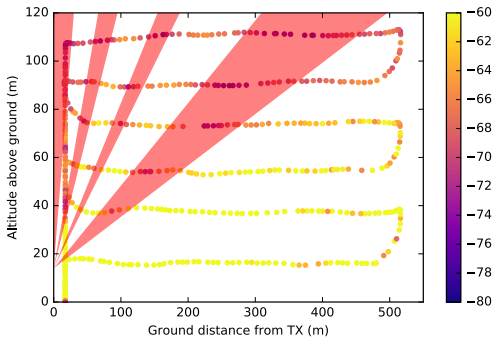



Fig. 17: Signal strength in dBm from a sector antenna. Nulls in signal strength are marked in red.



(a) Signal strength, 110 m AGL. (b) HG913P vertical gain pattern, 0 at 13 dBi.

Fig. 18: Sector antenna data. Nulls and antenna pattern correspond.

VI. APPLYING THE MODEL

Using the compositional model we have just developed, we will now examine several additional test cases, and compare our model with measured data. As a reference point, we will use a commercial-grade ray-tracer³, with the environment reconstructed from overhead photogrammetry. This analysis will be divided into three scenarios:

- Flying above the height of all buildings in the complex.
- Landing in the presence of a tall obstructing building.
- Emerging from behind a building.

First, we examine paths flown near the Hotel, the tallest building in the village. The bulk of the Hotel is 10.8 meters from the ground to the edge of the parapet, with an additional 3.8 meters to the top of the penthouse. We will model the building as a single knife-edge at the location of the wall closest to the receiver, without the additional height of the penthouse. Figure 19 shows the geometry of the two flights around the Hotel.

A UAS is most likely to spend most of its time well above the heights of the surrounding buildings. We would expect that up there, propagation can mostly be predicted purely from the two-ray effect, much as in the open field case. Figure 20 shows data, along with the modeled and ray-tracing results.




Fig. 19: Simplified geometry of Hotel and associated flight paths.

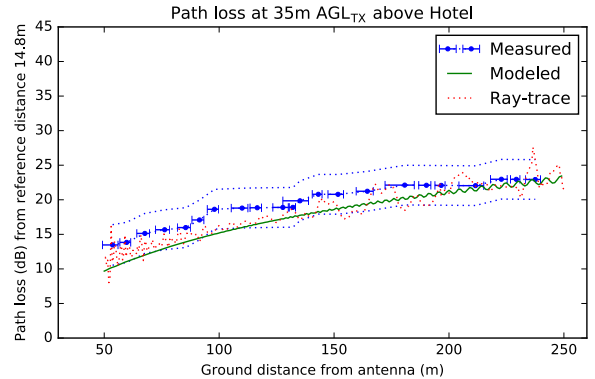


Fig. 20: Path loss along flight path 35m above base of Tx.

We note that, as expected, the data here matches closely with the open field scenario. The standard error $\sigma_{est} = 2.19$ dB for our model, and $\sigma_{est} = 1.94$ dB for the ray-traced simulation.

A UAS, while normally clear of buildings, will face a connectivity hazard as it comes closer to the ground, *e.g.* to deliver a package or get a closer look at a target. We flew a descent test starting at 50 meters AGL, shown in Figure 19. Figure 21 shows path loss as a function of altitude.

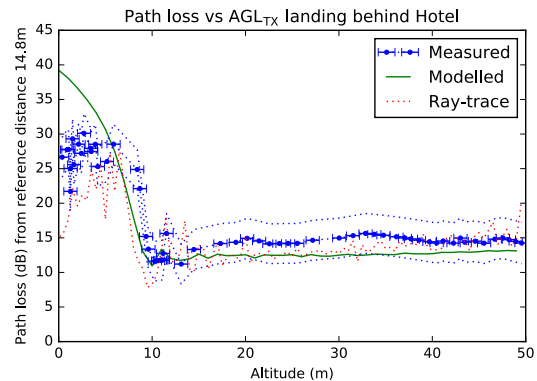


Fig. 21: Path loss vs. altitude behind the Hotel.

We see that again, the ray-tracing result is a bit better, with $\sigma_{est} = 4.43$ dB vs. $\sigma_{est} = 5.14$ dB for our compositional model. At the lower altitudes, where the building's obstruction is almost complete, our knife-edge approximation overestimates path loss as it does not account for diffraction

³Remcom Wireless InSite. <http://www.remcom.com/wireless-insite>

around the sides of the building; however, it still produces a useful upper bound on the path loss for avoidance purposes.




Fig. 22: Flight path of drone emerging from behind the House.

In the final scenario, the drone begins behind the House (Figure 12), where it is obstructed, and emerges to clear line-of-sight with the transmitter, shown in Figure 22. The drone maintains a low altitude well below the roof height of the House. Because the flight path is not perpendicular to the building, special consideration is needed when modeling the knife-edge. Along each Tx-Rx path, we map the knife-edge’s distance, if present, to the intersection between the path and the side of the building closest to the Rx, shown in red in Figure 22.

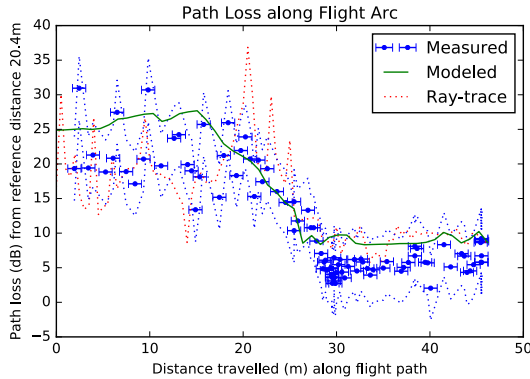


Fig. 23: Path loss along flight path. Note that the reference distance from which we computed path loss is different in this scenario.

In Figure 23, we see that the compositional model compares favorably in accuracy ($\sigma_{est} = 4.55$ dB) to the much more time-intensive ray-tracing solution ($\sigma_{est} = 4.86$ dB).

While 3D reconstruction and ray-tracing produces a more nuanced result, for these cases, a simple geometrical model based on the location of the nearest obstruction and distance to the transmitter produces a very comparable result, and should suffice for providing rapid estimations of path loss.

VII. MITIGATING HOLES-IN-THE-SKY

A. Simulating Coverage Holes

Let us explore how just the identified effects could create coverage holes for drones. Note that we do *not* claim that these are the only possible sources of path loss, only that they

are sufficient to create nulls in signal strength. We make the following assumptions for simulation:

- Obstacles are much wider than the width of the 1st Fresnel zone, and only diffraction over the top is considered.
- The noise floor is at least -100 dBm and < 10 dB of SINR [16] is poor signal, *i.e.* a coverage hole.
- Each cell site is operating on a different frequency channel, and propagation behaves as if the wavelength $\lambda = 0.33$ m.

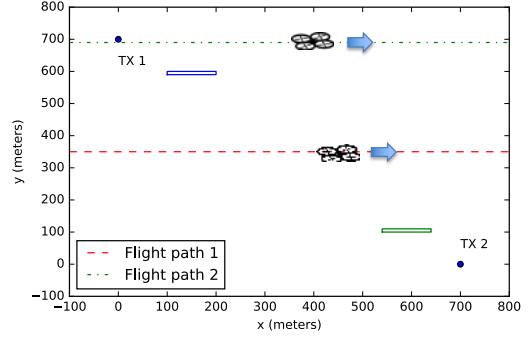


Fig. 24: Simulation geometry. Dotted lines show simulated flights, boxes show obstacles.

Figure 24 shows a simple scenario with two cell sites, operating at 10 W each. Tx 1 is 15 m off the ground, and Tx 2 is 22 m high. The sector antennas are tilted down by 15° . Let us take a look at flight path 1, at 30 m AGL. While this is a clear path in terms of obstacles, from a signal strength perspective, the buildings block the signal enough to create a coverage hole. Drones would not benefit from the micro- and pico-cells placed inside and around buildings to mitigate this problem for terrestrial customers.

If obstacles can block signal, why don’t we simply fly as high as possible to stay clear of any potential obstacles? High altitude brings other hazards—namely, the directional antennas have low gain and nulls directly above. Flight path 2 would run right into one of these antenna nulls.




Fig. 25: Coverage holes in 3-space from Fresnel zone incursion (center) and antenna nulls (sides), which appear as three-dimensional hazards for the UAS.

We plot regions where the signal strength is less than -90 dBm in Figure 25. We see that to a drone, a coverage map—and coverage dead zones—is a *three*-dimensional concept, and avoiding bad coverage could be achieved by navigation both above and around potential holes.

B. Detecting Coverage Holes

How do we ensure seamless wide-area coverage for UAS? One approach would be to augment current wide-area networks, using transmitters that direct RF energy into the air for high-traffic areas. Coverage maps pre-determined through modeling and testing as well as crowd-sourcing [9] would be used to inform drone flight planning.

Drones, however, will not operate solely on dedicated, statically-deployed, and pre-tested networks. Emergency services, disaster response, and military applications often require operation in communications-austere environments from rapidly-deployed transmitters. For these applications, the UAS will have to estimate, in real-time, the channel conditions it will face and take evasive action if necessary. Performing full 3D reconstruction of the scene and ray-tracing is a time-intensive process. A simple propagation model that only relies on knowledge of the nearest knife-edge could provide sufficient accuracy for avoiding holes-in-the-sky.

In the near-term, terrain and structure-aware ground control systems⁴ could be augmented with RF propagation modeling to aid mission planning. For fully autonomous UAS, we can combine computer vision methods [1] to determine the location and geometry of nearby obstacles, already necessary for sense-and-avoid, with given locations of nearby transmitters. This makes using a propagation model in real-time feasible, informing the autopilot of potential coverage holes.

VIII. RELATED WORK AND DISCUSSION

As we mentioned in Section II, the vast majority of work done in propagation modeling has been directed at ground-based receivers at roughly human height, with either high (cellular) or low (land-mobile radio) transmitters. Aerial RF surveys of existing cellular networks at high altitudes have in the past been conducted from manned aircraft [12], with the intent of placing cellular calls while in flight. More recently, there has been work in characterizing the path loss [17] and performance [18] of Wi-Fi signals at heights relevant to small UAS systems. These authors evaluate the performance of the Wi-Fi system at various distances and UAV orientations, as well as present a simple logarithmic (1D) model for the data. In urban environments, there has been work to adapt the pre-existing urban propagation models to UAS applications [13], taking into account diffraction off buildings. This work focuses on the point-to-point operator case, where the UAS communicates with a single ground station at human heights below the surrounding rooftops. We make the assumption that drone fleets will eventually operate on wide-area networks.

IX. CONCLUSION

Wide-area networks can be a good communications solution for commercial UAS operating in fleets and beyond line-of-sight, but were not built with aerial clients in mind. In this paper, we demonstrate the potential for coverage holes-in-the-sky through a series of field measurements. A transmitter was

deployed in a both an open field and a mock village, emulating a cell site. The resultant coverage was mapped using a UAS instrument. From these measurements, we identify two-ray ground interference, Fresnel zone incursion, and antenna nulls as the primary causes of coverage holes in our tested scenarios. We show that a deterministic model constructed from first principles can match reality with reasonable accuracy.

X. ACKNOWLEDGEMENTS

This work was partially funded by the Northrop Grumman Cybersecurity Research Consortium (NGCRC), and we thank our sponsor for their support. We would also like to thank Gerald Scott, Raymond Jackson, Tristan Allen, and the rest of the JIFX staff for their continued support and access to the JIFX event, without which these studies would not be possible.

REFERENCES

- [1] LSD-SLAM: Large-Scale Direct monocular SLAM. In *European Conference on Computer Vision*, 2014.
- [2] M. Boban, W. Viriyasitavat, and O. Tonguz. Modeling Vehicle-to-Vehicle Line of Sight Channels and its Impact on Application-Level Performance Metrics. Technical report, Carnegie Mellon University, Pittsburgh, PA, 2013.
- [3] D. Cichon and T. Kurner. Propagation prediction models. In *COST 231 Final Report*, pages 116–208, 1995.
- [4] R. Edwards and J. Durkin. Computer prediction of service areas for V.H.F. mobile radio networks. *Proceedings of the Institution of Electrical Engineers*, 116(9):1493, 1969.
- [5] ETSI. GSM 03.30 version 5.0.0, 1996.
- [6] M. Hata. Empirical Formula for Propagation Loss in Land Mobile Radio Services. *IEEE Transactions on Vehicular Technology*, 29:317–325, 1980.
- [7] J. Huang, F. Qian, A. Gerber, and Z. Mao. A Close Examination of Performance and Power Characteristics of 4G LTE Networks. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, pages 225–238, 2012.
- [8] B. Iannucci, P. Tague, O. Mengshoel, and J. Lohn. CROSSMobile: A Cross-Layer Architecture for Next-Generation Wireless Systems. *Silicon Valley Technical Report Series CMU-SV-14-001*, 2014.
- [9] J. Johansson, W. Hapsari, S. Kelley, and G. Bodog. Minimization of drive tests in 3GPP release 11. *IEEE Communications Magazine*, 50(11):36–43, 2012.
- [10] F. Pérez Fontán and P. Mariño Espiñeira. *Modelling the Wireless Propagation Channel: A Simulation Approach with MATLAB*. 2008.
- [11] T. S. Rappaport. *Wireless Communications: Principles and Practice*. 2002.
- [12] J. A. Romo, G. Aranguren, J. Bilbao, I. n. Odriozola, J. Gómez, and L. Serrano. GSM / GPRS Signal Strength Measurements in aircraft flights under 3,000 meters of altitude. *WSEAS Transactions on Signal Processing*, 5(6):219–228, 2009.
- [13] M. Simunek, P. Pechac, and F. P. Fontan. Excess loss model for low elevation links in urban areas for UAVs. *Radioengineering*, 20(3):561–560, 2011.
- [14] C. Sommer, S. Joerer, and F. Dressler. On the applicability of Two-Ray path loss models for vehicular network simulation. In *IEEE Vehicular Networking Conference, VNC*, pages 64–69, 2012.
- [15] E. Teng, J. D. Falcao, C. R. Dominguez, F. Mokaya, P. Zhang, and B. Iannucci. Aerial Sensing and Characterization of Three-Dimensional RF Fields. In *Second International Workshop on Robotic Sensor Networks*, Seattle, WA, 2015.
- [16] M. P. Wylie-Green and T. Svensson. Throughput, capacity, handover and latency performance in a 3GPP LTE FDD field trial. In *GLOBECOM - IEEE Global Telecommunications Conference*, 2010.
- [17] E. Yanmaz, R. Kuschnig, and C. Bettstetter. Channel measurements over 802.11a-based UAV-to-ground links. In *2011 IEEE GLOBECOM Workshops, GC Wkshps 2011*, pages 1280–1284, 2011.
- [18] E. Yanmaz, R. Kuschnig, and C. Bettstetter. Achieving air-ground communications in 802.11 networks with three-dimensional aerial mobility. In *Proc. IEEE INFOCOM*, pages 120–124, 2013.

⁴UgCS. <https://www.ugcs.com/>