6D Self-Localization of Drones using a Single Millimeter-Wave Backscatter Anchor

Maisy Lam[1,*], Laura Dodds[1,*], Aline Eid[2], Jimmy Hester[3], Fadel Adib[1,4]

[1]Massachusetts Institute of Technology, [2]University of Michigan, [3]Atheraxon, [4]Cartesian Systems *These authors contributed equally

mllam@mit.edu, ldodds@mit.edu, alineeid@umich.edu, jimmy.hester@atheraxon.com, fadel@mit.edu

***Abstract*—We present the design, implementation, and evaluation of MiFly, a self-localization system for autonomous drones that works across indoor and outdoor environments, including low-visibility, dark, and GPS-denied settings.**

MiFly performs 6DoF self-localization by leveraging a *single millimeter-wave (mmWave) anchor* in its vicinity - even if that anchor is visually occluded. MiFly's core contribution is in its joint design of a mmWave anchor and localization algorithm. The low-power anchor features a novel dual-polarization dual-modulation architecture, which enables single-shot 3D localization. MmWave radars mounted on the drone perform 3D localization relative to the anchor and fuse this data with the drone's internal inertial measurement unit (IMU) to estimate its 6DoF trajectory.

We implemented and evaluated MiFly on a DJI drone. We collected over 6,600 localization estimates across different trajectory patterns and demonstrate a median localization error of 7 cm and a 90th percentile less than 15 cm, even in low-light conditions and when the anchor is fully occluded (visually) from the drone. Demo video: youtu.be/LfXfZ26tEok

***Index Terms*—millimeter-wave, localization, drone**

I. Introduction

Recent years have seen growing interest in commercial drones for applications in indoor mapping [44], urban communications [8], autonomous delivery [4], entertainment [26], and search-and-rescue [45]. To enable such applications, these drones must localize themselves in their operational environments. In outdoor environments, drones typically rely on GPS to self-localize. However, GPS fails in indoor environments and urban canyons where satellites are obstructed by buildings. This has prompted researchers to explore other modalities for self-localization including vision/LiDAR, such as visual-inertial odometry (VIO) [41], [48]. These approaches have advanced the field of GPS-less localization, but often still struggle in dark and/or featureless environments (e.g., room with plain walls) [9]. Moreover, many of these vision-based approaches are more suitable for tracking *changes* in location rather than estimating the *absolute location*, which is critical for tasks like docking and delivery [4], [20].[1]

To address these challenges, past research in the wireless community has explored the use of Radio Frequency (RF) signals for indoor drone localization [27], [30], [40]. RF signals work in the dark and through fog, snow, and various materials (walls, cardboard, plastic, fabric, etc.) [17], which enables using them to localize in environments where visible light cannot work. In typical RF-based localization systems, multiple wireless anchors (e.g., UWB beacons) are placed in the environment, and a wireless transceiver on the drone localizes itself using the anchors. A key drawback of these approaches is the need for deploying numerous anchors in the environment for trilateration. This introduces significant overhead in cost and time to set up and maintain these anchors. Moreover, most existing systems require active and power-hungry anchors, which require tethering to a power source, further increasing the deployment requirements [2]. Thus, while state-of-the-art wireless localization systems for drones have demonstrated impressive accuracy (10-20 cm), they still suffer from significant infrastructure and maintenance requirements.

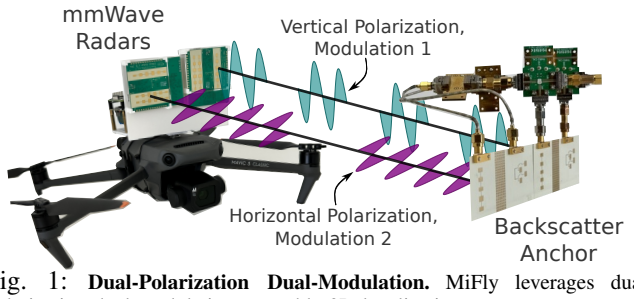

Fig. 1: **Dual-Polarization Dual-Modulation.** MiFly leverages dual-polarization dual-modulation to enable 3D localization.

This paper asks the following question: *Can we enable accurate 6DoF drone self-localization using only a single low-power wireless anchor that can be attached to a wall (similar to a sticker)?* Doing so would significantly reduce the infrastructure overhead of wireless localization because it would require a significantly smaller number of anchors, each of which is very easy to deploy. Moreover, low-power anchors can last for a long time on a small (coin-cell) battery.

In this paper we present MiFly, the first drone self-localization system capable of accurate (sub-decimeter) localization in dark, non-line-of-sight, and GPS-denied environments using only a single anchor. MiFly leverages millimeter-wave (mmWave) signals, which are used in modern radars and certain 5G systems. Specifically, the drone self-localizes by relying on 24 GHz radars mounted on the drone and a single custom-designed mmwave backscatter anchor placed in its vicinity.[2] By fusing the radar measurements with the drone's internal inertial measurement unit (IMU), MiFly estimates its 6DoF trajectory.

As shown in Fig. 1, MiFly operates by sending 24 GHz signals from an onboard radar to a planar anchor placed in the environment. MiFly uses 24 GHz mmWave radars because of their operational range and performance through occlusions as well as their compact form factor. [3] To operate at low-power, MiFly's anchor does not generate its own signal but

[1]One way to overcome this is to pre-map these environments, which introduces additional overhead and new challenges.

[2]Additional anchors can increase coverage area but are not required for localization.

rather modulates the reflections of the radar's signal back to the drone; i.e., it backscatters the mmWave radar signals. The signal reflected back from the anchor and received by the radar allows the drone to compute its location relative to the anchor.

Enabling fast and accurate localization using this approach is challenging for multiple reasons. The first challenge stems from estimating 3D location using commercial off-the-shelf 24 GHz mmWave radars. Specifically, existing 24GHz radars only provide antenna diversity in one dimension, limiting them to only estimate azimuth and range (i.e., 2D location). However, to support drone flight, localization must be three-dimensional. In principle, one could leverage two orthogonal mounted radars on the drone. This would allow us to measure both azimuth and elevation to the backscatter tag, and derive a 3D estimate. However, this design would suffer from interference between the radars. Furthermore, while radar measurements could be performed in series (e.g., measuring azimuth and elevation sequentially), this would introduce both latency and errors and prevent instantaneous 3D self-localization.

To enable high-speed localization, MiFly introduces a new method that synthesizes two techniques to decouple the interference between the two radars: polarization diversity and frequency diversity. Specifically, we co-design a *dual-polarization and dual-modulation* architecture in both our tag hardware and localization algorithm, as shown in Fig. 1. Here, the vertically polarized component of the tag backscatters the vertically polarized signal, while the horizontally polarized portion backscatters the horizontally polarized radar signal. Doing so simultaneously enables single-shot 3D localization. Moreover, the two orthogonally polarized components are designed to operate with different modulation frequencies, which enables the radars to isolate their reflections in the frequency domain. Not only does this minimize interference, but it also enables dealing with tilts that naturally occur during flight, which may otherwise cause shifts in polarization. We describe this method in detail in §III-A and how MiFly's joint algorithm-anchor design enables fast and efficient localization of the drone with respect to the anchor.

The second challenge is in converting the 3D anchor localization to drone self-localization (in the 3D world-frame). While the anchor measurements provide 3D estimates of the anchor relative to the radars, these estimates cannot be directly converted to the drone's position due to ambiguity caused by the drone's 6DoF (degrees of freedom) motion during flight. For example, if the drone rotates within the plane of an angular measurement, even without changing its location, the angle-of-arrival measurements of the tag will change and produce an incorrect position estimate (see §III-B for more details). In principle, to overcome this location ambiguity one could add more anchors to the environment and/or radars to the drone; however, such an approach is sub-optimal since it would increase deployment cost/overhead and drone payload.

[3]By comparison higher frequency radars (e.g., 60/77 GHz) suffer from shorter range, higher power consumption, higher cost, and more complex tag designs, while lower frequencies (e.g., RFIDs at 900 MHz) suffer from large antennas and low bandwidth.

To overcome the challenge, MiFly leverages the drone's IMU. While leveraging the IMU alone for localization would incur significant drift over time [49], MiFly only relies on the gyroscope to estimate the drone's rotation.[4] Not only does this approach disambiguate the 3D location, but it also enables 6DoF drone localization. We describe this ambiguity problem, and our method for mmWave-IMU fusion in more detail in §III-B.

We implemented MiFly on a DJI Mavic 3 Classic drone [14] using two Infineon Position2Go radars [22]. We custom-designed and fabricated 24 GHz backscatter anchors with orthogonally polarized patch antenna arrays on Rogers RO4350 substrate. Our evaluation was performed in multiple indoor environments, where we collected over 6,600 self-localization measurements during flight. Our evaluation demonstrates:

- MiFly achieves median errors of 4.8 cm, 1.0 cm, and 3.0 cm in the x, y, and z dimensions, respectively.
- MiFly performs almost equally well in line-of-sight (LOS) and non-line-of-sight (NLOS) scenarios, with a marginal drop in NLOS ($\sim$1 cm primarily due to SNR attenuation).
- MiFly's dual-polarization, dual-modulation design achieves $>$15 dB of isolation, allowing it to run radars simultaneously without interference and achieve low-latency 3D localization.

Contributions: We present MiFly, the first drone self-localization system capable of accurate (sub-decimeter) indoor 3D positioning using a single mmwave anchor. We introduce multiple contributions. First, we jointly design a mmWave backscatter anchor and localization scheme that leverage a dual-polarization and dual-modulation architecture to enable single-shot 3D localization. Second, we introduce a framework to combine mmWave and IMU data to achieve accurate 6DoF in-flight localization. This paper also contributes a end-to-end prototype implementation and real-world evaluation of the MiFly, demonstrating its accuracy in self-localizing in GPS-denied, indoor, dark and NLOS settings.

II. Related Work

Past work has explored different methods for GPS-less drone self-localization including vision, LiDAR, IMUs, wireless, and sensor fusion [23], [25], [35], [39]. Among these, the most ubiquitous are vision and lidar based systems; however, these struggle in dark, featureless environments. The most relevant to MiFly is past work that leverages Radio Frequency (RF) signals for drone localization, since these can operate in dark, GPS-denied, and low-feature environments.

The vast majority of past work on RF localization that is relevant to drones relies on placing wireless anchors in the environment and localizing with respect to these anchors. Researchers have investigated different wireless technologies for these anchors including WiFi [30] and Ultra-wideband [13], [37], [38], [40], [51]. Most of these works require deploying a dense network of anchors in the environment (e.g., dozens) to ensure that enough (at least 3) anchors remain within radio range for localization [30].

To reduce the infrastructure overhead, more recent work has

[4]Gyroscopes have demonstrated minimal ($<$1°/hr) drift for high performance grades [36].

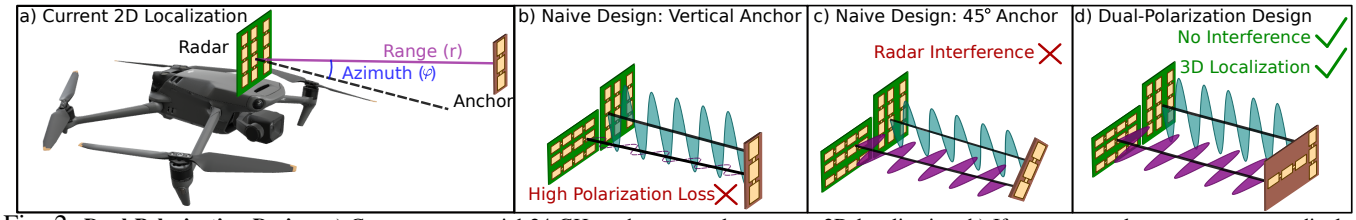

Fig. 2: **Dual-Polarization Design.** a) Current commercial 24 GHz radars can only measure 2D localization. b) If one were to leverage two perpendicularly mounted radars with a vertically polarized anchor, one radar would experience high polarization loss. c) Instead, if one were to leverage a 45° anchor, there would be radar interference. d) Finally, we introduce a new dual-polarization anchor design which can reflect both polarizations without interference.

investigated the potential to localize using a single anchor [10], [50], [52]. However, these systems still suffer from relatively large errors of 30-90 cm in 3D location estimation when leveraging a single anchor (WiFi AP); thus, they can be used for coarsely localizing a drone from an access point but cannot achieve higher level of accuracy typically required for drone self-localization and navigation. MiFly shares the goal of these systems but is the first to achieve sub-decimeter self-localization accuracy ($5-9\times$ higher accuracy than past systems that use a single anchor), which is important for applications that require higher resolution localization accuracy to support precise drone flight.

Research has also explored other mechanisms for RF localization that does not rely on anchors. Past work in this space falls in two main categories. The first leverages radars that track the drone from a distance, either from the vantage point of a base station [24], [34], [42], [46], [53] or from another tracker drone [16], [31]. These works that track from a distance struggle to uniquely identify drones and thus scale to multi-drone scenarios. This is due to the fact that they rely on Doppler and the motions of the drone body or propellers to separate the drone from the static background. Furthermore, they require sending data to the drone for localization, making them less reactive and introducing higher latency.

The second category has explored RF-based odometry to estimate change in position over time. Past work has included learning based approaches to fuse radar point clouds with IMU data [28], as well as doppler-based radar velocity estimation [28], [43] These approaches track the relative motion of the drone rather than absolute localization and have similar problems to VIO in that they may drift over time.

Finally, MiFly is related to past work on mmWave backscatter technology [5]–[7], [29], [47]. Recent proposals have investigated modulated mmWave backscatter for localization and have achieved accurate 2D localization leveraging a single anchor, but suffer from the same challenges of WiFi/UWB localization in terms of infrastructure overhead (number of anchors) to enable 3D localization [5], [7], [29], [47]. Other work for 6DoF localization [6], however, requires an array of multiple tags, and relies on heavy, power-hungry, and expensive ($3.5K) 77 GHz radars that would be difficult to mount on a drone. Finally, prior work using passive reflectors (e.g., without modulation) [21] have been shown, but require a large tag to produce necessary radar cross section and is susceptible to strong multi-path reflectors. MiFly builds on this nascent body of literature and is the first to enable not only 3D drone self-localization, but also 6DoF localization, using a single mmWave backscatter anchor. Moreover, it does so by leveraging commercially available off-the-shelf mmWave radars, designing a new anchor that addresses practical challenges with these radars, and fusing mmWave and IMU data to recover 6DoF poses.

III. Approach

A. Dual-Polarization Dual-Modulation 3D Localization

In this section, we first describe the challenges with 3D localization using commodity 24 GHz radars. Then, we describe MiFly's techniques to overcome them.

The Interference Problem: MiFly uses commodity 24 GHz radars to perform localization, as shown in Fig. 2a. The radar transmits a signal, which is modulated by the backscatter anchor, and received back at the radar. This signal can then be used to estimate the range (r), shown in purple, and the angle-of-arrival (φ), shown in blue. By combining the range and angle, the radar can be used to estimate the drone's 2D location. However, numerous indoor self-localization and navigation tasks require the drone to estimate its 3D location, which cannot be directly obtained from commercial off-the-shelf (COTS) 24 GHz radars.[5]

To overcome this challenge and obtain a 3D location, MiFly leverages two orthogonally mounted radars as shown in Fig. 1. In principle, the vertically mounted radar can estimate azimuth, the horizontally mounted radar can estimate elevation, and both can estimate the range. However, simply adding an additional orthogonal radar is insufficient to solve the 3D location estimation problem. This challenge in doing so is demonstrated in Fig. 2b. When mounting the radars orthogonally, the vertically mounted radar will have a vertical polarization (VP), shown by the teal sine wave, and the horizontally mounted radar will have a horizontal polarization (HP), shown by the purple sine wave. The anchor, which is vertically polarized, will only reflect the signal transmitted by the vertically polarized radar. The signal from the HP radar will be significantly attenuated, due to the high polarization mismatch between the HP radar and the VP anchor

Next, let us mathematically formalize this concept. One can define the polarization using a 2D complex vector $[E_h, E_v]$, where E_h and E_v are the complex-valued horizontal and vertical numbers applied to the transmitted signal [15]. A linearly polarized signal at angle θ would be defined as $[\cos(\theta), \sin(\theta)]$. Using this definition, we can find the signal received at any antenna by projecting the receiver polarization onto the transmitter polarization. Thus, the signal transmitted

[5]Recall COTS 24 GHz radars have antenna diversity in only one dimension.

by the HP radar and received by a VP anchor is:

$$\langle E_H, E^*_{anchor} \rangle = [1, 0] * \begin{bmatrix} \cos(\frac{\pi}{2}) \\ \sin(\frac{\pi}{2}) \end{bmatrix} = 0 \quad (1)$$

where E_H is the polarization of the HP radar, and E_{anchor} is the polarization of a anchor. This demonstrates polarization loss due to mismatch between the HP radar and VP anchor.

To address this loss, one approach might be to place the anchor at a 45° angle or leverage an anchor polarization that can receive and backscatter both radar signals (i.e., circular polarization). This approach can be seen in Fig. 2c. While this allows for both radars to receive reflections from the anchor, this approach leads to interference between the two reflected signals. This is because the anchor will backscatter both transmissions with a linear polarization at 45°, causing the backscattered signals to interfere with each other and making it difficult to separate them for localization.

Next, we formulate this interference problem. Without loss of generality, let us quantify the interference at the HP radar, or the signal transmitted from the VP radar, backscattered by a 45° anchor, and received at the HP radar. First, the signal transmitted by the VP radar and received by the anchor is:

$$\langle E_V, E^*_{anchor,45} \rangle = [0, 1] * \begin{bmatrix} \cos(\frac{\pi}{4}) \\ \sin(\frac{\pi}{4}) \end{bmatrix} = \sin\left(\frac{\pi}{4}\right) \quad (2)$$

where E_V is the polarization of the VP radar and $E_{anchor,45}$ is the polarization of a 45° anchor Next, the signal re-emitted from the anchor and received back at the horizontal radar is:

$$\langle E_{anchor,45}, E^*_H \rangle = [\cos(\tfrac{\pi}{4}), \sin(\tfrac{\pi}{4})] * \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \cos\left(\frac{\pi}{4}\right). \quad (3)$$

We can now quantify the interference signal at the HP radar:

$$\langle E_V, E^*_{anchor,45} \rangle \langle E_{anchor,45}, E^*_H \rangle = \sin\left(\frac{\pi}{4}\right) \cos\left(\frac{\pi}{4}\right) = 0.5.$$

Thus, the strength of the leakage is equal to 0.5. Next, we compare this to the strength of the received signal needed for localization. Following a similar derivation to above, the signal from HP radar to the anchor and back can be expressed as:

$$\langle E_H, E^*_{anchor,45} \rangle \langle E_{anchor,45}, E^*_H \rangle = \cos^2\left(\frac{\pi}{4}\right) = 0.5. \quad (4)$$

Here, we see that the strength of the leakage is actually equal to the strength of the localization signal (0.5). This demonstrates that using a 45° backscatter anchor would lead to significant interference when the radars are run in parallel. A similar derivation for an alternative circularly polarized anchor design (or designs that can reflect both HP and VP) would demonstrate that they incur the same interference problem.

While in principle, one could instead operate the two radars in series, this would incur additional localization latency, especially if the radar leverages multiple chirps for each location estimate. Similarly, adding coding on top of the chirps to disentangle them would also introduce additional latency, and make real-time localization more challenging.

This above derivation demonstrates that simply using two orthogonal radars is not sufficient to enable 3D localization.

Dual-Polarization Design: To overcome this challenge, MiFly introduces dual-polarization into the anchor design and localization algorithm. Our anchor design enables isolating the signals received by and reflected back to the HP and VP radars. To do this, we design our anchor with distinct sets of antennas, shown in Fig. 2d. The first set of antennas is horizontally polarized, and reflects only the chirps from the HP radar. The second set of antennas is VP and only reflects chirps from the VP radar. With this design, we are able to transmit from both radars simultaneously and minimize interference.

We can evaluate this design mathematically. The signals arriving at each anchor antenna from the HP radar are:

$$\begin{aligned} \langle E_H, E^*_{anchor,H} \rangle &= [1, 0] * \begin{bmatrix} 1 \\ 0 \end{bmatrix} = 1, \\ \langle E_H, E^*_{anchor,V} \rangle &= [1, 0] * \begin{bmatrix} 0 \\ 1 \end{bmatrix} = 0 \end{aligned} \quad (5)$$

where $E_{anchor,H}$ and $E_{anchor,V}$ are the polarization of the horizontally and vertically polarized antennas on the anchor, respectively. In this design, the anchor's horizontal antennas will receive (and reflect) all of the HP radar signal, while the anchor's vertical antenna will receive none of it. Therefore, the signal received back at the HP radar will be:

$$\langle E_H, E^*_{anchor,H} \rangle \langle E_{anchor,H}, E^*_H \rangle = 1$$

This shows that the HP radar receives back all of the power which it transmitted (without accounting for path loss, antenna gains, etc). Again, we can compare this to the interference received at the HP radar from the VP radar. In this case, the the signal transmitted by the VP radar, re-emitted by the VP anchor and received at the HP radar is:

$$\begin{aligned} \langle E_V, E^*_{anchor,V} \rangle \langle E_{anchor,V}, E^*_H \rangle &= \left([0, 1] * \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right) \left([0, 1] * \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right) \\ &= 0. \end{aligned}$$

Therefore, the strength of the interference signal is 0, while the strength of the desired signal for localization is 1. A similar derivation holds for the alternate case (i.e., from HP to VP).

This demonstrates the effectiveness of our dual-polarization design in reducing interference between the reflected signals.

Dual-Frequency Modulation: While in principle, polarization diversity provides perfect isolation between the two radar signals, this does not hold in practice. The backscattered signals will suffer from leakage, whereby the signal transmitted from the HP radar and reflected by the VP antennas on the dual-polarized anchor will be detected by the HP radar. This phenomenon is exacerbated in the case of a drone. The drone's rotations during flight may cause misalignment between the radars and the anchor, causing each radar to receive even more interference signal from the cross-polarized anchor antenna.

To understand this further, we conducted a real-world experiment to measure the leakage. Without loss of generality, we focused on measuring the leakage between the HP radar and the VP anchor antennas. Specifically, we transmit and receive on the HP radar, and only activate the VP anchor antennas to measure how much of the HP signals they would reflect.

In this experiment, the radars are placed approximately 1 m from the anchor, which is programmed to modulate signals at 350 kHz (only the VP antennas are programmed to modulate). We also apply 4° degree tilt (in roll) to the drone to simulate minor rotations during flight that can intensify leakage.

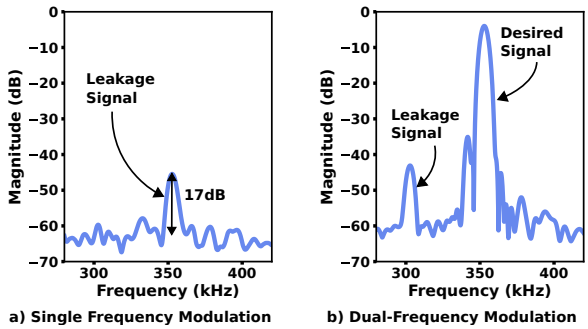

Fig. 3: **Dual-Frequency Modulation.** a) When transmitting a single frequency, there is 17 dB of interference. b) Our dual-frequency modulation shifts the leakage signal away from the desired signal in the frequency domain.

Fig. 3a. plots the magnitude of the signal received by the HP radar versus frequency. If the isolation were perfect, one would expect to see no peak in the received signal. However, as seen in the figure, there is a peak around 350 kHz, which is approximately 17 dB higher than the signal (or noise) baseline. This demonstrates the leakage present, despite the dual-polarization design.

In MiFly, we overcome this challenge through an additional technique: *dual-frequency modulation*. In this technique, we leverage different frequencies to modulate the vertical and horizontal antennas on the anchor, as illustrated in Fig. 1. By applying different modulation frequencies, we shift the signals for each polarization in the frequency domain for improved isolation. Specifically, the responses are separated in the frequency domain by:

$$\Delta f_{mod} = f_{mod,H} - f_{mod,V} \tag{6}$$

where $f_{mod,H}$, and $f_{mod,V}$ are the modulation frequencies of the HP and VP antennas. Using this technique, we can see that Δf_{mod} allows for separation of responses from the desired polarization and the leakage of the orthogonal polarization.

We conduct a real world experiment to validate our dual-frequency modulation technique. We use the same setup as previously described and measure the received signal at the HP radar. However, we now apply our dual-frequency technique to modulate both the HP and VP antennas on the anchor. Specifically, we shift the modulation frequency of the VP antennas to 300kHz, and modulate the HP antennas at 350kHz.

Fig. 3 plots this result. The plot now displays a peak of 17 dB, but it is located near 300 kHz. This is the leakage signal that has been shifted by the dual-frequency modulation. There is also a ∼50 dB tall peak at 350 kHz. In comparison, this peak is the desired localization signal (from the HP radar to the HP anchor). Now, the leakage and desired signals can be easily separated. This shows the benefit of our dual-frequency modulation for improving isolation between polarized signals.[6]

Finally, one might wonder whether dual-modulation would have been sufficient for 3D localization, without the need for dual-polarization. However, that would not be sufficient or desirable. First, we still need dual-polarization on the radar transmitters themselves (to obtain 3D locations). Second, the dual-polarized tag achieves a 6 dB gain over a simple linearly-polarized one (e.g., at 45°) due to better polarization matching.

Single-Shot 3D Localization: Now that the two signals are easily separable, they can be used for localization. Since we are using FMCW radars, the backscatter peak frequency shown in Fig. 3b will be directly proportional to both the modulation frequency and the anchor's range. It can be defined as:

$$f_{backscatter} = f_{mod} + f_{range} \tag{7}$$

where $f_{backscatter}$ is the peak frequency of the backscatter's response (for the corresponding polarization), f_{mod} is the corresponding modulation frequency, and f_{range} is the frequency shift corresponding to the anchor's range.

Therefore, we can use this peak frequency to isolate the backscatter response and perform 3D localization. We describe this localization at a high level, and refer readers to [29] for more details. Our localization follows three steps:

1) First, to find the backscatter frequency, we take an FFT of the received signal. We then search for a peak frequency in a window near the corresponding modulation frequency as:

$$\hat{f}_{backscatter} = \underset{f \in (f_{mod}, f_{mod}+W)}{\arg\max} |FFT_{RX}(f)| \tag{8}$$

where $\hat{f}_{backscatter}$ is the estimated backscatter frequency, W is the window size, and $FFT_{RX}(f)$ is the FFT of the received signal evaluated at frequency f.

2) Second, we compute the range to the anchor as [29][7]:

$$\hat{r}_{anchor} = \frac{c(\hat{f}_{backscatter} - f_{mod})}{2k} \tag{9}$$

where $\hat{r}_{anchor}$ is the estimated range to the anchor, c is the speed of light, and k is the slope of the FMCW chirp signal.

3) Finally, we can compute the angle-of-arrival (AoA) to each radar, using the phase difference between the two RX antennas. Specifically, we take the difference between the phases of the FFT peaks between each RX signal ($\Delta\phi$):

$$\Delta\phi = \angle FFT_{RX,1}(\hat{f}_{backscatter}) - \angle FFT_{RX,2}(\hat{f}_{backscatter}) \tag{10}$$

where $FFT_{RX,1}$ and $FFT_{RX,2}$ are the FFTs of the received signals from RX antenna 1 & 2. Then, we compute AoA [29]:

$$\varphi = \sin^{-1}\left(\frac{\Delta\phi\lambda}{2\pi d}\right) \tag{11}$$

where λ is the wavelength of the signal and d is the spacing between the two RX antennas.

We repeat the above three steps for both radars, giving us the range, azimuth, and elevation to the anchor [8] Then, our final 3D anchor location is defined as:

$$x = \hat{r}_{anchor} * \sin(\varphi) * cos(\theta), \quad y = \hat{r}_{anchor} * \cos(\varphi) * cos(\theta)$$
$$z = \hat{r}_{anchor} * \sin(\theta)$$

where x, y, and z are the coordinates of the anchor location, φ and θ are the azimuth and elevation estimates, respectively.

Finally, it is worth noting that we employ an outlier rejection

[6]We note that some small amount of signal could travel from the VP radar to the HP anchor to the HP receivers, causing interference. However, between the polarization diversity and the dual-frequency modulation, the interference is significantly reduced, such that remaining interference will not be strong enough to hinder the localization.

[7]In practice, the anchor's modulation produces two peaks in the FFT. We use these two peaks to compute the anchor's frequency $\hat{f}_{backscatter}$ without relying on a known value of f_{mod}. We refer readers to [29] for more details.

[8]We average the range estimates from our two radars to improve accuracy.

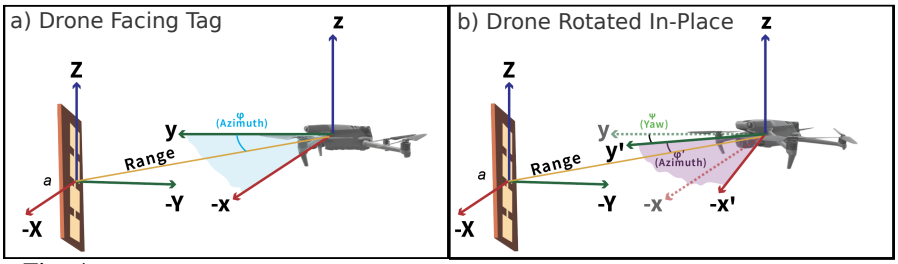

Fig. 4: **Impact of Heading** a) Drone directly facing anchor. b) Drone rotated about the Z axis.

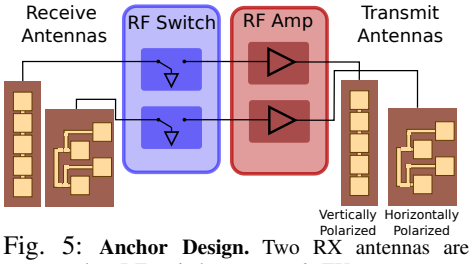

Fig. 5: **Anchor Design.** Two RX antennas are connected to RF switches, amps, & TX antennas.

mechanism to remove low quality measurements. First, we define the signal-to-noise ratio (SNR) of a measurement as:

$$SNR = \frac{||FFT(f_{peak})||^2}{\sum_{f \in \{f_L, f_H\} - \{f_{peak}\}} ||FFT(f)||^2} \tag{12}$$

where $FFT(f)$ is the FFT at frequency f, $||\cdot||^2$ denotes the magnitude squared (power), and f_L, f_H are the bounds of the search space over which we search for the anchor's response[9]

Given this SNR, we can remove any outlier measurements with a SNR below a threshold P_{thresh}[10]:

$$SNR < P_{thresh}. \tag{13}$$

B. RF-IMU Fusion

The previous sections described how we produce a high-speed 3D location from a single anchor. In this section, we first describe the challenge with converting these 3D anchor estimates to 3D drone self-localization estimates. Then, we describe how MiFly overcomes this challenge and additionally produces 6DoF drone pose estimates.

As described in §III-A, our mmWave localization technique produces a single-shot 3D location of our anchor. One might think that these estimates can simply be inverted to provide the drone's 3D location. However, these anchor-based 3D estimates cannot be used directly as the position of the drone during flight. This is because the drone's rotational movement during flight causes ambiguity by entangling information from the drone's position and rotation in the AoA measurement. Thus, any rotational movement can lead to error in the estimated 3D location of the drone.

We illustrate an example of this in Fig. 4. In Fig. 4a, the drone's radars are directed towards the plane of the anchor and the drones frame during flight is aligned with the anchors frame a. In this case, the drone is at a distance of 2 m in x and 2 m in y from the anchor. Therefore, the radar would estimate an azimuth of 45°, as shown by the blue azimuth angle (φ). In Fig. 4b, the drone (and its radars) are at the same 3D location, but have rotated 20° about the yaw, as shown by the green yaw angle (ψ). In this case, the azimuth detected by the radars will only be 25° (φ' in purple). If we were to directly use this azimuth estimate to produce a 3D drone location, this would results in localization error of 80 cm in x and 56 cm in y. This error is even further exacerbated as the range to the anchor increases, since the azimuth offsets produce location errors which are scaled by the distance. Additionally, this problem is further complicated during simultaneous rotation in both yaw and pitch, as ambiguity in the location is a complex function of 6DoF pose.

To overcome this ambiguity, our system fuses the drone's internal IMU with the mmWave localization to account for the variations in heading. To enable self-localization, our goal is to solve for the position of the drone in the world frame (p^w_{drone}), where we define the world frame w as global frame of reference. To start, recall from §III-A that we obtain an estimate of anchor's location from the radars, defined as:

$$p_a^{drone} = (x, y, z) \tag{14}$$

where p_a^{drone} is the 3D position of the anchor in the drone's frame of reference during flight. Additionally, from the internal IMU, we can determine the rotation of drone's frame during flight with respect to the world frame R^w_{drone} :

$$R^w_{drone} = R_X(\alpha)R_Y(\beta)R_Z(\gamma) \tag{15}$$

where R_X, R_Y, R_Z are the elementary rotation matrices. α, β and γ are the IMU euler angles.[11] Using this rotation and the mmWave-based anchor location, we then apply a transformation to determine the desired location of the drone in the world frame. Formally:

$$p^w_{drone} = -R^w_{drone} * p_a^{drone} + t_a^w \tag{16}$$

where t_a^w is the known transformation between the global reference frame and the anchor location.

Thus, the final 6DoF pose estimate of the drone during flight with respect to the world is $(p^w_{drone}, R^w_{drone})$.

IV. Implementation & Evaluation

mmWave Anchor Design. A real world photo of the anchor is shown in Fig. 1, and the anchor design is shown in Fig. 5. We custom-fabricated our anchor design on Rogers RO4350 substrate. The anchor consists of two receive antennas, one horizontally polarized and one vertically polarized. The signals received at these antennas are then fed through HMC547ALC3 SPDT switches which routes the signals to a matched load, or feeds them into HMC342LC4 amplifiers. After the amplifiers, the signals are routed to an identical set of antennas for re-emission. Our anchor currently consumes 260 mW.[12]

Physical Setup. We implement MiFly on a DJI Mavic 3 Classic [14]. We mount two Infineon Position2Go 24 GHz radars [22], a Raspberry Pi 4 Model B and UPS power supply on the drone with custom 3D printed parts. We connect the radars to the Pi via USB. We note that, in the future, the radar could be powered entirely off the drone's battery and

[9]We remove a small gap around f_{peak} from the sum due to peak width.

[10]We additionally filter out measurements that have a large difference in estimated range between the two receiver antennas on the radar.

[11]We align the IMU reference orientation to align with the world frame.

[12]This is significantly lower power than approaches operate with a single anchor (e.g., 5-25W per WiFi anchor [11]). It is also lower than most multi-anchor approaches, even on a per-anchor basis [2].

Axis	10^{th} pctl	Median	90^{th} pctl
X	0.9 cm	4.8 cm	11.8 cm
Y	0.2 cm	1.0 cm	3.2 cm
Z	0.5 cm	3.0 cm	9.4 cm
Roll	0.1°	0.5°	1.2°
Pitch	0.5°	1.3°	2.4°
Yaw	0.4°	1.3°	3.1°
3D	**2.8 cm**	**7.2 cm**	**14.2 cm**

TABLE I: **Localization & Rotation Error.** The 10^{th}, 50^{th}, 90^{th} pctl errors in X, Y, Z, roll, pitch, yaw, & 3D.

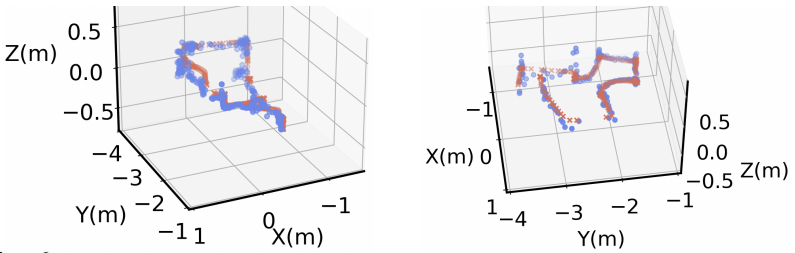

Fig. 6: **Flight Path.** Two example trajectories with MiFly's measurements (blue) and the interpolated ground truth measurements (red) plotted in 3D.

only reduce the overall battery life by ∼1 minute (from 46 minutes to 45 minutes).[13] The total payload is approximately 385 grams. We transmit chirps of 300μs duration over 200 MHz of bandwidth and average all 16 chirps within a frame to produce a single localization measurement. We perform a one-time calibration and apply it to all measurements to correct for offsets in the radar and tag hardware.

Software. We combine the IMU information and RF measurements in post processing on a laptop running Ubuntu 22.04. Since we do not have access to the drone's software to synchronize the IMU with the Raspberry Pi, we perform a calibration to measure the time offset between the two.[14] After applying the time offset, we then interpolate the IMU data to provide heading estimates at the times that the radar measurements were captured. Since the two radars are not perfectly clock synchronized, we pair the closest two measurements in time to provide a single 3D estimate.

Evaluation Environment. We evaluate MiFly in different environments: a typical office setting, with tables, chairs, cabinets, etc. (e.g., high multipath), and a flight space for further range and angle evaluation. We note that our experiments were conducted with people/objects moving in the background, as well as typical wireless interference (WiFi, 5G, Bluetooth, etc).

Ground Truth. To measure the localization accuracy of our system, we use the OptiTrack™ Motion Capture System to provide 6DoF tracking. We place tracking markers on the drone and continuously extract the 6DoF pose of the drone during flight. In a case that the drone is temporarily missed by the OptiTrack system, we remove that section of the trajectory from the evaluation. Since the ground truth is recorded on a separate computer than MiFly, we use the same NTP server to provide time synchronization. We then interpolate the ground truth measurements to find the ground truth at the times that MiFly's measurements were recorded.

Metric. We measure the localization error as the 3D L2 norm difference between MiFly's predicted location and the ground truth location. We also measure the rotation error as the difference in degrees between the drone's estimated euler angles and the ground truth euler angles.

Baseline. We implement a visual inertial odometry (VIO) baseline for comparison.[15] We leverage an Intel® RealSense™ Tracking Camera T265. We mount the camera on the drone using 3D printed mounts and use the cameras internal tracking algorithm to extract the 3D coordinates of the drone.

[13]The Position2Go radar only consumes 2.1W [22], while the drone averages ∼100W during flight [14]. Therefore, powering the radars from the drone's battery would have a negligible impact on the drone's battery life.

[14]We rotate the drone at the start of a trajectory and correlate the change in the IMU's yaw with the change in radar's azimuth to find the time offset.

V. Performance Results

A. 6DoF Error

To evaluate MiFly's overall 6DoF error, we conducted experiments where the drone was flown in the environment and MiFly continuously tracked its location. Outlier measurements were filtered using the procedure described in §III-A. The IMU data was time aligned with mmWave measurements and used to derive the final 6DoF estimates (§III-B). The ground truth was simultaneously tracked (§IV). We collected over 6,600 localization estimates (after filtering) across 31 separate trajectories. We collected measurements in both visual line-of-sight (LOS) and non-line-of-sight (NLOS) of the tag[16], in light and in the dark, and while flying 1D, 2D, and 3D trajectories.

Table. I shows the 10$^{\text{th}}$, 50$^{\text{th}}$, and 90$^{\text{th}}$ percentile localization errors in X, Y, Z, roll, pitch, yaw, and 3D L2 norm. We note:

- MiFly achieves a median L2 norm 3D error of 7.2 cm and a 90$^{\text{th}}$ percentile less than 15 cm. Additionally, the 90$^{\text{th}}$ percentile errors in roll, pitch, and yaw are 1.2°, 2.4°, and 3.1°, respectively. This demonstrates MiFly's ability to accurately self-localize using a single anchor during flight.
- Interestingly, the median error in Y is 1 cm, while the median in X and Z are 4.8 cm and 3.0 cm respectively. This is due to the fact that error in Y is primarily determined by error in our range estimate, while error in X and Z are primarily determined by error in our azimuth and elevation estimates. In general, MiFly has slightly higher accuracy in range than AoA estimates. However, the overall 3D error still remains low enough to enable drone self-localization.

B. Qualitative Results

Next, we show two qualitative results. Fig 6a-b show two example trajectories with MiFly's measurements (blue) and the interpolated ground truth measurements (red), plotted in 3D. MiFly successfully reconstructs the trajectory, with the estimated locations closely matching the ground truth.

C. Comparison to VIO

Next, we compare the performance of MiFly to that of our VIO baseline (See §IV). We define two types of environments. The first are well-lit, feature-rich environments favorable to VIO. The second have low-light and fewer features, presenting a challenge for VIO systems. We run three different trajectories in both of these environments.

[15]We compare to VIO as it is primarily used for indoor drone self-localization [32]. While recent work has began exploring LIDAR [18], these sensors are expensive ($2k), heavy for smaller indoor drone deployments, and suffer in poor visibility (fog, smoke, etc) & featureless environments [3], [19].

[16]We use LOS/NLOS to refer to visual LOS/NLOS.

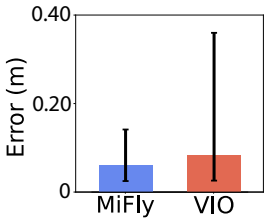

Fig. 7: **Low Visibility/Low Feature Environments.** The 10^{th}, 50^{th}, 90^{th} pctl 3D errors in low light/featureless environments for MiFly and VIO.

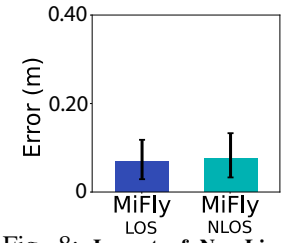

Fig. 8: **Impact of Non-Line-of-Sight on Accuracy** The 10^{th}, 50^{th}, and 90^{th} pctl of the 3D errors of MiFly in LOS&NLOS.

Fig. 7 shows the median 3D error of MiFly and VIO in low-light, low feature environments. The error bars denote the 10th and 90th percentiles. We note that MiFly is able to achieve a 90th percentile error of 14 cm. On the other hand, VIO is unable to track with a high level of accuracy, with a 90th percentile of 36 cm. We note that these errors are due to complete loss of tracking. Over longer trajectories and without pre-mapping of the environment, we estimate the VIO performance would degrade even further. This shows the need for MiFly's mmWave localization, that can operate in low-light, feature-less environments.[17]

VI. Microbenchmarks

A. LOS vs NLOS

In our first microbenchmark, we evaluate the impact of visual line-of-sight on MiFly's localization accuracy. We flew the drone along specific trajectories by following markers on the floor. We repeat each trajectory twice, once when the anchor was in LOS of the drone and once when the anchor was hidden behind clutter such that it was in visual non-line-of-sight of the drone throughout the entire trajectory. We ran four different trajectories for a total of 8 trajectories and over 1,200 localization measurements for both LOS and NLOS.

Fig. 8 shows the median 3D error of MiFly in LOS and NLOS. The error bars show the 10th and 90th percentiles. We note that at all percentiles, the localization error in NLOS is less than 1 cm larger than in LOS. This demonstrates that MiFly is able to successfully self-localize even in practical environments that include clutter and occlusions.

B. Error vs Range

Next, we evaluate the error of MiFly as a function of the distance to the anchor. We flew the drone along a straight line directly in front of the anchor, starting roughly 1 m away and ending about 8 m away. We grouped MiFly's measurements into different bins based on how far the drone was from the tag at the time of the measurement. We measured the 3D localization error with and without our outlier rejection described in §III-A. We also measured the frame retention rate (FRR), or the percentage of measurements within a given distance bin not removed by our outlier rejection.

Fig. 9 shows the 3D error in meters as a function of the distance to the anchor. The bar graph shows the median error both with (red) and without (blue) outlier rejection, with the error bars denoting the 10th and 90th percentile.[18] The dotted lines show the FRR. We make the following remarks:

- MiFly has a median and 90th percentile filtered 3D error less than 14 cm and 25 cm up to 5 m, respectively. [19]
- Beyond 5m, the 90th percentile of the unfiltered range error increases above 50cm. This is mainly due to limitations of MiFly's hardware (e.g.,antenna array size, sampling rate, transmit power etc.), which prevents it from successfully localizing at these ranges. Various improvements can be made to further boost performance in the future (see §VII). Further, the outlier rejection removes these measurements to avoid large errors.

C. Error vs Rotation

Next, we evaluate MiFly's error as a function of the drone's yaw rotation. We flew the drone at a fixed location, and rotated 360 degrees in both directions. We grouped measurements into bins based on the drone's yaw at the time of the measurement (where 0 yaw is facing the tag). Since yaw rotation primarily impacts the azimuth estimation, we evaluated azimuth error as a function of the drone's rotation.

Fig. 10a shows the error in azimuth (degrees) as a function of the drone's yaw angle (degrees). The bar graph shows the median error both with (red) and without (blue) outlier rejection, with the error bars denoting the 10th and 90th percentile. The dotted lines show the FRR. We note:

- Between -50° and 40°, the median filtered error remains below 5.1°, while the FRR is 100%. This demonstrates that MiFly is able to operate at a wide range of rotations.
- At extreme angles of -70° and 70°, the median unfiltered error increases to 127° and 109°. This is due to the fact that the radar's antenna design has a limited field-of-view (FoV), and cannot accurately measure the tag's response at this angle. This demonstrates the need for MiFly's outlier rejection to remove these unreliable measurements.

D. Error vs Elevation

Next, we evaluate the error of MiFly as a function of the drone's elevation to test performance at different heights. We flew the drone in a straight line vertically. We grouped measurements into different bins based on the drone's elevation relative to the tag at the time of the measurement. Since the relative height to the tag primarily impacts MiFly's elevation estimate, we evaluated the error in the elevation.

Fig. 10b shows the error in elevation (degrees) as a function of the drone's ground truth elevation (degrees). The bar graph shows the median error both with (red) and without (blue) outlier rejection, with the error bars denoting the 10th and 90th percentile. The dotted lines show the FRRs. We note:

- Between -30° and 20°, the median filtered error remains below 1.4° and the FRR remains above 75%. This shows that MiFly is able to estimate elevation at a wide range of heights.
- At an extreme elevation angle of -40°, the median unfiltered error increases to 25°. This is again due to the limited radar FoV, and shows the benefit of MiFly's outlier rejection.

E. Radar Isolation

Our final microbenchmark measures the isolation provided by MiFly's dual-polarization dual-modulation design. To do

[17] In well-lit environments, the VIO tracks with a median of 2.1 cm and a 90th percentile of 4.1 cm, as expected.

[18] We remove any bars from the bar plot with 5 or fewer measurements, since they are not statistically significant.

[19] Note that since this microbenchmark is only a small subset of our flights designed for evaluating the impact of distance on our accuracy, these numbers are not expected to match the median & 90th reported in §V-A.

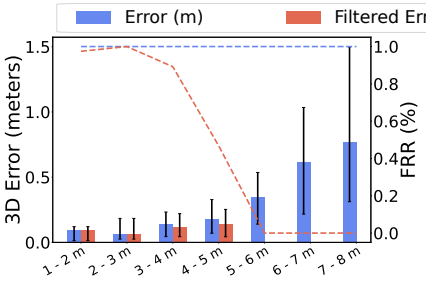

Fig. 9: **3D Error vs Distance.** Plot of median 3D error & FRR vs distance with (red) and without (blue) filtering. Error bars denote 10th&90th pcntl.

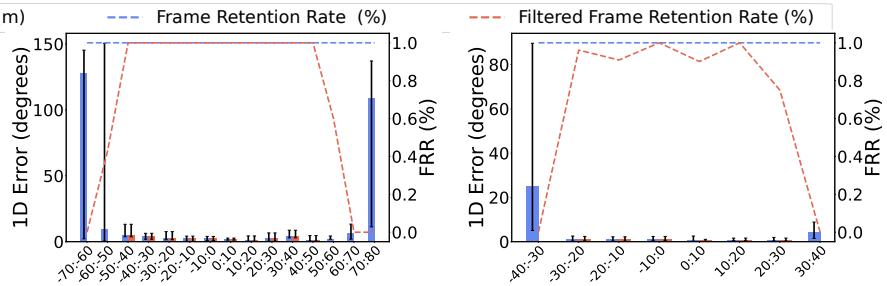

(a) Azimuth Error vs Rotation. (b) Elevation Error vs Elevation.

Fig. 10: **Angular Microbenchmarks.** (a) plots the error & FRR (dashed lines) in azuimuth vs the drone's yaw angle with (red) and without (blue) filtering. (b) Error in elevation vs the drone's elevation.

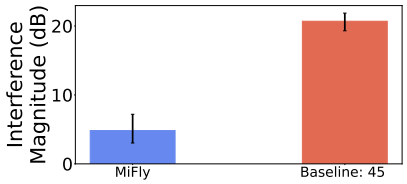

Fig. 11: **Anchor Isolation.** The isolation between the radars when leveraging MiFly's anchor design (blue) and a 45° anchor (red).

so, we compare the leakage between the two radars when using our design compared to a baseline 45° linearly-polarized tag.

We placed the drone and its radars 0.7 m from the anchor. We programmed the anchor to modulate its VP antennas at 300 kHz and its HP antennas at 350 kHz. To measure the leakage, we transmitted from the VP radar and received from the HP radar. To ensure the measurement contained only the interference between the radars and not signal transmitted from the HP radar, we set the HP radar parameters to minimize transmitted power and place a metal reflector to block the TX antennas. We measured the interference as the magnitude of the FFT peak corresponding to the interfering signal (relative to the surrounding noise). We then repeated this experiment with a 45° tag modulated at 350 kHz. We collected over 45 experimental trials for each anchor design.

Fig. 11 plots the interference magnitude in dB for MiFly (blue) and a 45° tag (red). The bar graph shows the median error, with error bars denoting the 25th and 75th percentiles. We note that MiFly experiences a median interference of 4.8 dB, while the 45° tag experiences median of 20.7 dB. Therefore, MiFly's dual-frequency dual-modulation design provides 15.9 dB of isolation, showing our technique's value.

VII. Discussion and Limitations

Occlusions: MiFly does not need to be in visual line-of-sight of the anchor to operate; however, it requires mmWave signals to traverse occlusions on the direct path between the anchor and the drone. MmWave signals can traverse everyday occlusions (such as walls [17]), but are significantly attenuated by certain materials (such as reinforced concrete & metal). In cases of complete blockage, MiFly would temporarily lose accurate tracking until the drone moved to a location where the direct path was no longer fully blocked. Further, additional anchors could be deployed to limit cases of complete blockage.

Coverage: We demonstrated MiFly's robust performance up to 5 m in §VI-B. While this is a promising start, there are two important things to note. First, MiFly could extend the operational range of an anchor to hundreds of meters via power amplifiers, improved radars, improved antenna designs, etc. [33], [47]. The angular field-of-view can also be increased with the use of omnidirectional antennas. Second, in scenarios where very large areas need to be covered (e.g., an urban city) additional anchors can be deployed so that the drone can switch between different anchors as it travels.

Extension to Multiple Drones: There are many ways to extend these techniques to operate with multiple drones, such as transmitting different (orthogonal) waveforms from different drones for multiple access to the same anchor.

Anchor Design: The ability to leverage a single anchor for localization significantly reduces the infrastructure overhead for indoor wireless drone self-localization. MiFly's anchor design can be made even more compact and easier to deploy. In particular, its planar structure (which we fabricated on a printed circuit board) can be printed as a sticker similar to RFIDs, which are also backscatter technologies. Moreover, the design can be transitioned to more optimized antenna structures (e.g., retroreflective) that can make it more efficient, potentially eliminating the need for power amplifiers altogether and making the anchor entirely passive.

Incorporating Doppler: At high speeds, it is possible for Doppler shifts to impact localization. However, commercial deployment of indoor drones for applications such as warehouse inventorying are typically limited to maximum flight speeds of 2-3 mph [1], [12]. At these speeds, Doppler shifts would induce a maximum error of 5 cm, allowing MiFly to achieve accurate localization. That said, to increase the accuracy during faster flights, it would be interesting future work to explore choosing range from the range-Doppler domain.

VIII. Conclusion

We present MiFly, a self-localization system for autonomous drones using only a single millimeter-Wave anchor. MiFly leverages a novel dual-polarized, dual-modulated mmWave anchor and mmWave-IMU Fusion self-localization algorithm to ultimately achieve precise, high speed 6D localization. Our experiments verify that MiFly can robustly self-localize in GPS denied environments, adverse lighting conditions, and in non-line-of-sight to the anchor. For future work, we aim to incorporate our self-localization into the navigation of the drone and enable autonomous flight for applications such as docking, delivery and discovery.

Acknowledgments. We thank the anonymous reviewers and the Signal Kinetics group for their help and feedback. This research is sponsored by NSF (Award #2313234) and the MIT Media Lab.

REFERENCES

[1] Dexion. https://www.dexion.biz/product/scan-drones/.

[2] 95Power. Uwb indoor positioning system.

[3] EOS Data Analytics. Lidar vs. radar: Differences uses to pick the right one, 2024.

[4] Robin Riedel Andrea Cornell, Sarina Mahan. Commercial drone deliveries are demonstrating continued momentum in 2023, 2023.

[5] Kang Min Bae, Namjo Ahn, Yoon Chae, Parth Pathak, Sung-Min Sohn, and Song Min Kim. Omniscatter: extreme sensitivity mmwave backscattering using commodity fmcw radar. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services*, MobiSys '22, New York, NY, USA, 2022. ACM.

[6] Kang Min Bae, Hankyeol Moon, and Song Min Kim. Supersight: Sub-cm nlos localization for mmwave backscatter. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services*, MOBISYS '24, New York, NY, USA, 2024. ACM.

[7] Kang Min Bae, Hankyeol Moon, Sung-Min Sohn, and Song Min Kim. Hawkeye: Hectometer-range subcentimeter localization for large-scale mmwave backscatter. In *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services*, MobiSys '23, page 303–316, New York, NY, USA, 2023. ACM.

[8] Lorenzo Bertizzolo, Michele Polese, Leonardo Bonati, Abhimanyu Gosain, Michele Zorzi, and Tommaso Melodia. mmbac: Location-aided mmwave backhaul management for uav-based aerial cells. In *Proceedings of the 3rd ACM Workshop on Millimeter-wave Networks and Sensing Systems*, pages 7–12, 2019.

[9] Carlos Campos, Richard Elvira, Juan J. Gómez Rodríguez, José M. M. Montiel, and Juan D. Tardós. Orb-slam3: An accurate open-source library for visual, visual–inertial, and multimap slam. *IEEE Transactions on Robotics*, 37(6):1874–1890, 2021.

[10] Guoxuan Chi, Zheng Yang, Jingao Xu, Chenshu Wu, Jialin Zhang, Jianzhe Liang, and Yunhao Liu. Wi-drone: Wi-fi-based 6-dof tracking for indoor drone flight control. ACM MobiSys '22, page 56–68.

[11] Cisco. Cisco aironet 3800 series access points.

[12] Corvus. Corvus warehouse inventory drone. https://www.corvus-robotics.com/corvus-one.

[13] Stéphane D'Alu, Oana Iova, Olivier Simonin, and Hervé Rivano. Demo: In-flight localisation of micro-uavs using ultra-wide band. In *Proceedings of the 2020 International Conference on Embedded Wireless Systems and Networks*, EWSN '20, USA, 2020. Junction Publishing.

[14] DJI. Dji mavic 3 classic. https://www.dji.com/mavic-3-classic.

[15] Laura Dodds, Isaac Perper, Aline Eid, and Fadel Adib. *A Handheld Fine-Grained RFID Localization System with Complex-Controlled Polarization*. ACM, New York, NY, USA, 2023.

[16] Sedat Dogru and Lino Marques. Pursuing drones with drones using millimeter wave radar. *IEEE RAS*, 5(3):4156–4163, 2020.

[17] Kairui Du, Ozgur Ozdemir, Fatih Erden, and Ismail Guvenc. Subterahertz and mmwave penetration loss measurements for indoor environments. In *2021 IEEE ICC Workshops*, pages 1–6, 2021.

[18] Flyability. Elios 3, 2023.

[19] Flyability. What is a lidar drone?, 2023.

[20] Carlo Giorgio Grlj, Nino Krznar, and Marko Pranjić. A decade of uav docking stations: A brief overview of mobile and fixed landing platforms. *Drones*, 6(1):17, 2022.

[21] Tatsuya Iizuka, Takuya Sasatani, Toru Nakamura, Naoko Kosaka, Masaki Hisada, and Yoshihiro Kawahara. *MilliSign: mmWave-Based Passive Signs for Guiding UAVs in Poor Visibility Conditions*. Association for Computing Machinery, New York, NY, USA, 2023.

[22] Infineon. Demo position2go. https://www.infineon.com/cms/en/product/evaluation-boards/demo-position2go/.

[23] Ninad Jadhav, Weiying Wang, Diana Zhang, Swarun Kumar, and Stephanie Gil. Toolbox release: A wifi-based relative bearing sensor for robotics. *arXiv preprint arXiv:2109.12205*, 2021.

[24] Jia Zhang, Xin Na, Rui Xi, Yimiao Sun, Yuan He. mmhawkeye: Passive uav detection with a cots mmwave radar. *arXiv preprint arXiv:2308.06479*, 2023.

[25] Swarun Kumar, Stephanie Gil, Dina Katabi, and Daniela Rus. Accurate indoor localization with zero start-up cost. MobiCom '14. ACM.

[26] Michael J. de la Merced Lauren Hirsch. Fireworks have a new competitor: Drones, 2023.

[27] Sangmin Lee, Seungho Yoo, Joon Yeop Lee, Seongjoon Park, and Hwangnam Kim. Drone positioning system using uwb sensing and out-of-band control. *IEEE Sensors Journal*, 22(6):5329–5343, 2022.

[28] Chris Xiaoxuan Lu, Muhamad Risqi U Saputra, Peijun Zhao, Yasin Almalioglu, Pedro PB De Gusmao, Changhao Chen, Ke Sun, Niki Trigoni, and Andrew Markham. milliego: single-chip mmwave radar aided egomotion estimation via deep sensor fusion. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, 2020.

[29] Charles A. Lynch, Ajibayo O. Adeyeye, J. G.D. Hester, and Manos M. Tentzeris. When a single chip becomes the rfid reader: An ultra-low-cost 60 ghz reader and mmid system for ultra-accurate 2d microlocalization. In *2021 IEEE International Conference on RFID (RFID)*.

[30] Rufo I. Marasigan, Yolanda D. Austria, Jennifer B. Enriquez, Luisito Lolong Lacatan, and Rhowel M. Dellosa. Unmanned aerial vehicle indoor navigation using wi-fi trilateration. In *11th IEEE ICSGRC*, 2020.

[31] Christos Milias, Rasmus B. Andersen, Bilal Muhammad, Jes T. B. Kristensen, Pavlos I. Lazaridis, Zaharias D. Zaharis, Albena Mihovska, and Dan D. S. Hermansen. Uas-borne radar for autonomous navigation and surveillance applications. *IEEE T-ITS*, 2023.

[32] ModalAL. Vio development drones.

[33] Siddharth Mohan, Kshitiz Bansal, Manideep Dunna, and Dinesh Bharadia. Retroreflection with amplification for long range mmwave sensing.

[34] Peter Joseph Basil Morris and K. V. S. Hari. Detection and localization of unmanned aircraft systems using millimeter-wave automotive radar sensors. *IEEE Sensors Letters*, 5(6):1–4, 2021.

[35] A. Moura, J. Antunes, A. Dias, A. Martins, and J. Almeida. Graph-slam approach for indoor uav localization in warehouse logistics applications. In *2021 IEEE ICARSC*, pages 4–11, 2021.

[36] Advanced Navigation. Inertial measurement unit (imu) – an introduction.

[37] Thien Minh Nguyen, Abdul Hanif Zaini, Kexin Guo, and Lihua Xie. An ultra-wideband-based multi-uav localization system in gps-denied environments. In *2016 International Micro Air Vehicles Conference*.

[38] Pat Pannuto, Benjamin Kempke, and Prabal Dutta. Slocalization: Sub-uw ultra wideband backscatter localization. IPSN '18. IEEE Press, 2018.

[39] Jie Qian, Kaiqi Chen, Qinying Chen, Yanhong Yang, Jianhua Zhang, and Shengyong Chen. Robust visual-lidar simultaneous localization and mapping system for uav. *IEEE Geoscience and Remote Sensing Letters*.

[40] Jorge Peña Queralta, Carmen Martínez Almansa, Fabrizio Schiano, Dario Floreano, and Tomi Westerlund. Uwb-based system for uav localization in gnss-denied environments: Characterization and dataset. In *2020 IEEE/RSJ IROS*, pages 4521–4528, 2020.

[41] Mohammad Fattahi Sani and Ghader Karimian. Automatic navigation and landing of an indoor ar. drone quadrotor using aruco marker and inertial sensors. In *2017 IConDA*, pages 102–107, 2017.

[42] Andrey Sesyuk, Stelios Ioannou, and Marios Raspopoulos. 3d millimeter-wave indoor localization. In *2023 13th IPIN*, 2023.

[43] Emerson Sie, Zikun Liu, and Deepak Vasisht. Batmobility: Towards flying without seeing for autonomous drones. In *Proceedings of the 29th Annual International MobiCom*, ACM MobiCom '23, New York, NY, USA, 2023. ACM.

[44] Ishveena Singh. Skydio drones can now scan indoor spaces autonomously.

[45] Skydio. How drones are used for search and rescue, 2023.

[46] Suhare Solaiman, Emad Alsuwat, and Rajwa Alharthi. Simultaneous tracking and recognizing drone targets with millimeter-wave radar and convolutional neural network. *Applied System Innovation*, 6(4), 2023.

[47] Elahe Soltanaghaei, Akarsh Prabhakara, Artur Balanuta, Matthew Anderson, Jan M. Rabaey, Swarun Kumar, and Anthony Rowe. Millimetro: Mmwave retro-reflective tags for accurate, long range localization. MobiCom '21, page 69–82. ACM, 2021.

[48] Amr Suleiman, Zhengdong Zhang, Luca Carlone, Sertac Karaman, and Vivienne Sze. Navion: A 2-mw fully integrated real-time visual-inertial odometry accelerator for autonomous navigation of nano drones. *IEEE Journal of Solid-State Circuits*, 54(4):1106–1119, 2019.

[49] Dr.-Ing. Edgar v. Hinüber. If you intend to use an inertial measurement system. https://www.imar-navigation.de/downloads/Decision_assistant-Dateien/Decision_assistant.pdf.

[50] Deepak Vasisht, Swarun Kumar, and Dina Katabi. Decimeter-level localization with a single wifi access point. NSDI'16. USENIX Association.

[51] S Zahran, MM Mostafa, A Masiero, AM Moussa, A Vettore, and N El-Sheimy. Micro-radar and uwb aided uav navigation in gnss denied environment. *The ISPRS-Archives*, 42:469–476, 2018.

[52] Shengkai Zhang, Wei Wang, and Tao Jiang. Wi-fi-inertial indoor pose estimation for microaerial vehicles. *IEEE TIE*, 2020.

[53] Peijun Zhao, Chris Xiaoxuan Lu, Bing Wang, Niki Trigoni, and Andrew Markham. 3d motion capture of an unmodified drone with single-chip millimeter wave radar. In *IEEE ICRA*, 2021.