

RetroLiDAR: A Liquid-crystal Fiducial Marker System for High-fidelity Perception of Embodied AI

Kenuo Xu[¶], Bo Liang[¶], Jingyu Li[¶], Chenren Xu^{¶✉*}

[¶]School of Computer Science, Peking University

[✉]Key Laboratory of High Confidence Software Technologies, Ministry of Education (PKU)

ABSTRACT

As embodied AI gradually transitions into practical applications, enhancing the fidelity of how embodied agents perceive the physical world has become a critical challenge. Current perception methods typically rely on computer vision-based fiducial marker systems, which suffer from limitations such as insufficient reading distance, poor localization accuracy, and high susceptibility to environmental lighting conditions. Currently, SPAD sensor-based LiDAR technology is emerging in commercial mobile devices due to its compact size, high precision, and low power consumption. This paper presents the design of the RetroLiDAR system, which chimes with the concept of backscatter in wireless technology, to create a liquid-crystal fiducial marker system that can be directly read by LiDAR. On the marker side, we use retroreflective materials to reflect the LiDAR's emitted light back and employ a liquid crystal modulator to adjust the intensity of the light signal. On the LiDAR end, we design a signal processing pipeline to demodulate the marker's modulation message using the temporal received signal strength. Experimental results from our prototype demonstrate that compared to visual fiducial markers, RetroLiDAR extends the reading distance by 2.6x compared to QR codes and by 44% compared to AprilTags, while reducing the median ranging error by 85%. We also present a low-power marker circuit design, a link budget analysis, and two proof-of-concept applications to validate the system's efficacy and practicality.

CCS CONCEPTS

• **Hardware** → **Wireless devices**; • **Computer systems organization** → **Embedded systems**.

KEYWORDS

LiDAR; Fiducial Marker; Liquid Crystal; Visible Light Backscatter Communication

ACM Reference Format:

Kenuo Xu, Bo Liang, Jingyu Li, Chenren Xu. 2025. RetroLiDAR: A Liquid-crystal Fiducial Marker System for High-fidelity Perception of Embodied AI. In *The 23rd ACM Conference on Embedded Networked Sensor Systems (SenSys '25)*, May 6–9, 2025, Irvine, CA, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3715014.3722073>

*✉: chenren@pku.edu.cn

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SenSys '25, May 6–9, 2025, Irvine, CA, USA

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1479-5/2025/05...\$15.00

<https://doi.org/10.1145/3715014.3722073>

1 INTRODUCTION

Embodied AI is recognized as crucial for advancing towards Artificial General Intelligence (AGI) and serves as the foundation for a broad range of applications that seamlessly bridge cyberspace with the physical world. Unlike conventional (disembodied) AI, which excels at executing specific tasks within virtual environments, embodied AI is designed to navigate the complexity and unpredictability of the real world [1, 2]. This shift towards embodied systems emphasizes not just abstract computation, but also general-purpose perception and interaction through physical agents [3].

One of the most significant challenges in advancing embodied AI lies in the perception of objects and their spatial relationships within the physical world [4]. Unlike conventional object recognition in static images, embodied agents are required to perceive and respond to changes in dynamic, three-dimensional spaces, which entails a deeper understanding of spatial relationships and evolving contexts [5, 6]. To establish a robust foundation for embodied perception, we argue that *fidelity* is a critical metric that must be prioritized. In this context, fidelity refers to the system's ability to achieve three core objectives: (1) accurately identifying objects in 3D space, especially those lacking distinct visual features; (2) understanding and extracting spatial information about these objects from a distance, including their positions and relationships within a scene; and (3) maintaining high-accuracy perception under real-world changes, such as shifts in position, lighting variations, and interactions from other users. While it is possible to perform embodied perception using only the cameras on mobile devices in structured environments with strong visual features, these methods struggle when applied over longer distances, in environments lacking distinct features, or in challenging lighting conditions [7–9].

To complement camera-based systems, (AR) fiducial marker systems [10–12] are often used in (embodied) robotics [13–15]. Fiducial markers are typically patterned tags that can be detected and identified using computer vision algorithms, allowing for (1) positioning of the markers and (2) identification of the ID encoded within them. While fiducial markers represent a significant improvement over purely camera-based approaches, their current limitations include performance constraints, with decimeter-level accuracy typically achieved only when positioning markers (several centimeters in size) at distances of a few meters [16]. The identification and ranging performance degrade with smaller markers. The limitations of pure camera-based systems push us to leverage more precise sensors in order to perceive the physical world with higher fidelity.

Among the various sensors available, LiDAR has been the most accurate sensor to capture 3D content from the physical world and inputting it into the digital realm. Particularly, the new generation of solid-state LiDAR based on SPAD sensors to achieve direct Time-of-Flight (ToF) measurements offers millimeter-level

wide-angle depth information. These devices consume power only in the hundred-milliwatt range and have physical sizes measured in millimeters. The emerging technology has been recently integrated into new types of commercial mobile devices such as the Apple iPhone and Vision Pro [17, 18]. We believe that such available LiDAR sensors on mobile devices are the most high-fidelity technologies for embodied perception, especially for robots with size and power limitations.

However, despite LiDAR’s superiority in ranging accuracy, further research is needed to fully utilize LiDAR for embodied perception. LiDAR inherently lacks the ability to capture rich semantic and texture information, and its spatial resolution remains insufficient [19–22]. The LiDAR point cloud exhibits further degradation with increasing distance and is susceptible to spatial misalignment artifacts and environmental corruptions, resulting in compromised object detection performance (*e.g.*, error rates exceed 50% in practical scenarios) [23, 24]. Therefore, the distance maps (*i.e.*, common LiDAR point clouds) generated by LiDAR are not suited for place recognition with and without fiducial markers. In contrast, modern LiDAR systems often provide an *intensity map* for each point. A straightforward approach is to use this intensity map to capture black-and-white marker patterns by detecting variations in intensity [25–28]. However, the current limitations in LiDAR’s spatial resolution require markers to be disproportionately large. For example, a marker would need to be as large as 1.2 by 1.2 meters to be detected at a distance of 10 meters, even with a powerful LiDAR [25]. Such large marker sizes severely limit practical applications. The fundamental constraint of LiDAR here is still the spatial resolution compared with RGB cameras. Although researchers are exploring ways to integrate vision systems with LiDAR to enable collaborative operation [29–32], performance is still bottlenecked by the vision system: if identification fails due to lighting or distance issues, any advantage in ranging becomes irrelevant. Therefore, to fully utilize LiDAR for fiducial marker reading, it is crucial to overcome these limitations and develop a LiDAR-only solution.

In this paper, we pose the research question: can we leverage wireless technology to complement LiDAR, enabling it to read fiducial markers without compromising its distance measurement capabilities? Inspired by recent advancements in Visible Light Backscatter Communication (VLBC) [33], we design a liquid-crystal fiducial marker system that can be directly read by LiDAR. The basic principle of VLBC involves using electrical signals to modulate the transparency of liquid crystals, thereby encoding information by altering the intensity of incident light. We find that the same principle can be applied to the near-infrared light used by LiDAR. More importantly, this passive modulation aligns with the working principle of LiDAR, which is designed to receive only the light it emits, preserving its normal ranging capabilities. At the core of our marker design is a liquid crystal modulator (LCM). When the LCM is charged or discharged, its transparency changes, resulting in high or low intensity signals in the LiDAR’s intensity map. Additionally, we incorporate a retroreflector into the marker to significantly enhance the signal-to-noise ratio (SNR) received by the LiDAR. To decode the ID information encoded in the marker, we design a detection algorithm that detects the marker with high sensitivity and low false alarm rate. We also design a specialized receiving circuit that captures the high-frequency patterns of the LiDAR signal to

wake up the system to minimize overall power consumption. We further conduct a link budget analysis to examine the relationship between marker size, reading distance, and LiDAR configurations to guide the design of markers for practical scenarios.

We implement a prototype system of RetroLiDAR to validate our design. Our experiments demonstrate that using a small $3 \times 3 \text{ cm}^2$ marker and a low-power mobile LiDAR, our system effectively reads markers at a distance of 2.6 meters, with a median ranging error of 1.3 mm under various ambient light conditions. For a tiny $1 \times 1 \text{ cm}^2$ marker, the system can read it from a distance of 1.3 meters. Compared to conventional visual fiducial markers, RetroLiDAR extends the reading distance by 2.6x compared to QR codes and by 44% compared to AprilTags, while reducing the median ranging error by 85%. We also design two use cases of embodied AI in indoor environments to validate their feasibility through experiments.

Contributions.

- We introduce a novel fiducial marker concept utilizing liquid crystal modulators that can be read by LiDARs on mobile devices, offering long-range, high-precision, and robust performance against ambient light variations.
- We present the design of the RetroLiDAR system (§3), featuring a low-power marker hardware with dual filters and a full signal processing pipeline at the LiDAR, demonstrating the capabilities of using LiDAR for reading liquid-crystal fiducial markers.
- We develop a prototype (§5), and evaluation results (§6) show that RetroLiDAR outperforms visual fiducial markers in both reading distance and ranging precision. Experiments in use case studies (§7) validate its feasibility for embodied perception.

2 BACKGROUND

2.1 LiDAR for Embodied Perception

LiDAR (Light Detection and Ranging) is a sensing technology that measures the distance to objects by illuminating them with laser light and detecting the reflected pulses. The working principle of a LiDAR is to emit laser pulses and measure the time it takes for each pulse to travel to an object and reflect back to the sensor. The time measurement is then directly converted into distance based on the speed of light. This process, known as direct Time-of-Flight (dToF), provides high-accuracy measurements and is commonly applied in fields such as autonomous driving and robotics.

Although LiDARs have been regarded as powerful, expensive, and large devices, recent advances in Single-Photon Avalanche Diode (SPAD) sensors have enabled the development of compact, low-power, and eye-safe LiDAR systems. SPAD sensors feature very high sensitivity, allowing LiDAR to maintain accuracy with low power consumption. In addition, advancements in semiconductor manufacturing processes have made it possible to create SPAD sensor arrays, enabling LiDARs that require no mechanical scanning components. These SPAD-based LiDAR devices have been successfully commercialized for use in mobile devices to capture spatial data. For example, recent Apple devices such as the iPhone, iPad, and Vision Pro are equipped with SPAD LiDARs that demonstrate sub-cm-level ranging accuracy at distances up to 5 meters [34].

In SPAD-based LiDAR systems, each measurement cycle involves the emission of multiple laser pulses at a fixed frequency (typically 15-200 kHz [35]), where individual SPAD pixels (analogous to the

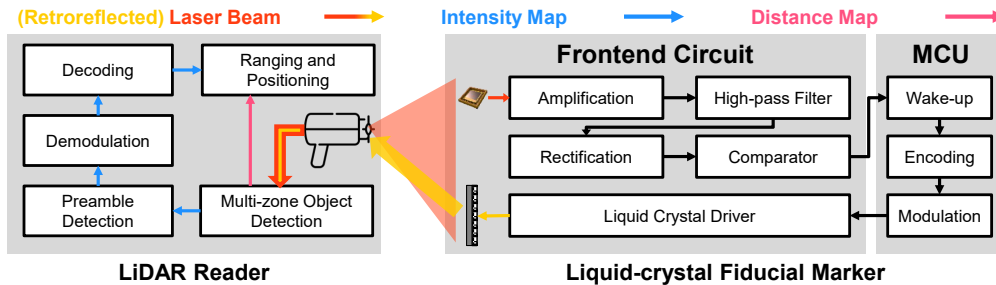


Figure 1: The architecture and workflow of RetroLiDAR.

pixels of a frame in a camera) sequentially accumulate ToF data to construct ToF histograms. This raw data output requires internal digital signal processing to be transformed into LiDAR ranging information. After processing, for each pixel (x, y) in a frame at time t , the output format is a quadruple (D, I, A, S) , where D represents the distance, I the received signal intensity, A the ambient light level, and S the status of the measurement (normal or not). This information is then provided to the user for further application.

2.2 Liquid Crystal Modulators

A liquid crystal modulator (LCM) is a thin, low-power, and cost-effective optical modulator that controls light transmission or blocking. Its working principle is similar to that of a TFT LCD display. Specifically, the core of the LCM is a layer of liquid crystal enclosed by two glass electrodes. In its natural state (*i.e.*, no voltage applied), the liquid crystal molecules rotate the incident light’s polarization by 90° due to intermolecular forces from internal structures. In contrast, a voltage (*e.g.*, 5 V) aligns the liquid crystals when applied to the electrodes, maintaining the original light polarization. The liquid crystal layer is sandwiched between two perpendicular polarizing filters. As a result, when no voltage is applied, the LCM remains highly transparent, whereas applying voltage reduces its transparency. The functionalities of LCM allow us to modulate temporal optical signals by applying a time-varying electrical signal to the LCM electrodes. It is worth noting that the LCM does not emit light itself, but rather *modulates any light passing through it*.

Leveraging the LCM’s ability to control the passage of light, several methods can be designed to represent binary bits. The most straightforward approach is on-off keying (OOK). In OOK, the presence of a carrier signal for a specific duration represents a binary one, while its absence represents a binary zero. Therefore, we can encode “0” when no voltage is applied to the LCM, resulting in higher transparency, and encode “1” when voltage is applied, reducing its transparency. The advantage of this approach (as opposed to reversing the representation of “0” and “1”) is that the liquid crystal remains in a high-transparency state when idle, which enhances the overall signal strength.

3 RETROLIDAR SYSTEM DESIGN

3.1 Design Overview

The basic design of RetroLiDAR is to leverage the modulation capability of liquid crystals for near-infrared light of LiDAR, as shown in Fig. 1. The system consists of a LiDAR reader and liquid-crystal markers. Each marker incorporates three major components: a photodiode and associated circuits for LiDAR signal reception, a LCM with the requisite control circuitry, and an MCU to manage the

overall system logic. The LiDAR reader operates as a standard dToF sensor, similar to conventional depth sensing systems.

During the operation of RetroLiDAR, the LiDAR continuously emits near-infrared laser pulses. These high-frequency signals are detected by a photodiode on the marker. After amplification, the signal is passed through a high-pass filter to selectively isolate the LiDAR signal. Following rectification, if the signal strength surpasses a predetermined threshold, the marker identifies the presence of the LiDAR signal, triggering a wake-up call to the MCU. The MCU then initiates the encoding and modulation of message bits. The modulated electrical signal is applied to the LCM, which modulates the intensity of the reflected laser back to the LiDAR.

On the LiDAR side, the raw received signal is first processed internally through a multi-zone object detection algorithm, which generates intensity and distance maps for all pixels within the sensor’s field of view (FoV). The intensity map contains the modulated temporal information from the marker. A preamble detection algorithm is employed to identify the pixels that contain active markers. These candidate pixels are subsequently demodulated and decoded, with successfully decoded markers (meeting error correction criteria) being paired with their corresponding distance measurements from the distance map. This fusion of decoded marker IDs and precise ranging data produces the final marker reading output.

3.2 Liquid-crystal Optical Frontend

Necessity of LCM. One straightforward approach to transmitting ID information to a LiDAR system is to adopt an active light source (*e.g.*, an infrared LED) to modulate the information that can be received by the SPAD sensor in the LiDAR. However, aside from the issue of high power consumption, this approach fundamentally conflicts with the working principle of LiDAR, which is designed to receive only the light it emits, treating all other light sources as noise. To investigate the impact of using an external LED, we conduct an experiment where a 940-nm LED (operating in the same frequency band as the LiDAR) is placed 1 meter away from the LiDAR at varying power levels. As shown in Tab. 1, the presence of the LED not only reduces the percentage of normal frames but also significantly affects the LiDAR’s ranging accuracy. Even at very low power (0.08 W), the standard deviation of multiple measurements increases by a factor of 5. This result demonstrates that an active




Figure 2: Signal intensity w/ and w/o retroreflector.

Table 1: Impact of LED on LiDAR performance.

LED Power (W)	0	0.08	0.25	0.40	0.70	0.83	1.30
Normal Frame Percentage	100	100	100	100	72	74	0
Std. Dev. of Ranging (mm)	1.91	9.58	10.27	11.20	19.86	17.89	NaN

light source can severely disrupt the normal operation of the LiDAR. Therefore, our marker needs to transmit information passively, which is why we choose the LCM.

Retroreflector. An important aspect of LiDAR’s internal signal processing is its tendency to select and output the strongest signal of each pixel. The incorporation of LCMs inherently introduces optical attenuation due to the cumulative light absorption by their multilayer components, including the front/rear polarizers due to Malus’s law. Therefore, we need to maximize reflectivity of our marker to enhance the signal received by the sensor. A key property that can be leveraged is retro-reflectivity. Unlike typical objects that scatter light diffusely, retroreflectors redirect incident LiDAR emissions back along their original path, resulting in more concentrated returns for the sensor to capture. We measure the signal intensity detected by the LiDAR from a $3 \times 3 \text{ cm}^2$ retroreflector compared to a non-retroreflective PCB board at various distances, as shown in Fig. 2. The intensity is measured in units of Kcps/SPAD (Kilo-count per second per SPAD), which quantifies the number of photons detected by the SPAD array and directly correlates with the light intensity. The results demonstrate an average $100.4\times$ enhancement in signal intensity when using retroreflectors. More importantly, we empirically observe that ambient objects reflect an average signal of 4 Kcps/SPAD. If the marker’s signal strength falls below this level, the LiDAR is unable to detect it. For markers without a retroreflector, this limitation occurs beyond 1.4 meters, whereas the retroreflector significantly extends this distance. These findings underscore the necessity of using a retroreflector. We will further analyze the properties of retroreflector in §4.

3.3 Marker Hardware Design

To maintain low power consumption, RetroLiDAR markers should operate predominantly in sleep mode, activating only when queried by a LiDAR system. Therefore, we need to design a circuit-level wake-up mechanism capable of selectively triggering the marker upon LiDAR detection. The straightforward approach to receiving LiDAR signals is to equip the marker with a photodiode (PD); however, it proves inadequate in practice. Mobile LiDAR systems typically adhere to the Class I laser transmitter standard, with power output below 5 mW. Moreover, they utilize laser emitters that produce highly divergent beams. Our testing indicates that when using only a PD as receiver, the LiDAR signal is overwhelmed by ambient noise at distances greater than 20 cm. Therefore, an effective wake-up mechanism must achieve both high sensitivity and selectivity for LiDAR signals. Our design addresses these requirements through the integration of two specialized filtering stages:

Spectrum Filter. LiDAR uses a laser with a wavelength of 940 nm in the near-infrared band. By attaching a 940-nm bandpass optical filter to the PD, we can reduce the PD’s response to ambient light by several orders of magnitude, as light at this wavelength constitutes only a small portion of ambient light. This filtering has almost no effect on the LiDAR signal itself. With the spectrum filter, we can use a high-gain amplifier, which can extend the work range without the risk of saturation from ambient light.

Signal Pattern Filter. The operational mode of LiDAR imparts a unique frequency signature to their emitted pulse signals, which directly corresponds to the pulse reception rate as detailed in §2.1.




Figure 3: An example of LiDAR pulse patterns.

Upon analyzing a segment of the signal as shown in Fig. 3, we identify a unique frequency characteristics: high-frequency components at 46.1 kHz and 54.2 kHz during each measurement. This pattern arises from the working principles of LiDAR, where multiple independent laser pulses are emitted in quick succession within a single measurement to accumulate sufficient photons (§2). Consequently, this high-frequency feature of laser pulses enables a high-pass signal filter that can selectively detect the LiDAR signal.




Figure 4: RetroLiDAR marker wake-up circuit.

Optical Reception Circuit Design. Fig. 4 illustrates our circuit design. We use a PD covered with a spectrum filter to receive the LiDAR signal. After passing through a transimpedance amplifier, the signal is filtered by a high-pass filter to retain only the high-frequency components. The signal is then amplified a second time and passed through a rectifier, converting it into a stable voltage signal that is highly correlated with the incident LiDAR signal. After comparing this signal with a properly set threshold, it can serve as the trigger for the wake-up circuit, which is subsequently processed by additional logic circuits.

3.4 Marker Data Transmission

Modulation and Coding. Visual fiducial markers typically encode information through spatial patterns, such as variations in brightness or contrast. In contrast, RetroLiDAR utilizes LCM that encodes information in the temporal domain. We use OOK modulation to represent binary bits with light signals (§2.2). To ensure robust marker ID transmission in the presence of random errors, a coding system is required to maximize the number of correctable bit errors while minimizing false identification/alarm rates. Specifically, our system adopts Binary Golay code [36], which encodes 12 bits of data into a 24-bit word with a Hamming distance of 8. This coding allows for the correction of up to 3-bit errors and the detection of up to 4-bit errors. Notably, the 12-bit information length is common for fiducial markers. For example, the widely used AprilTag tagStandard41h12 family supports 2,115 different markers, with a dictionary size smaller than 12 bits.

Preamble and Packet Format. Similar to visual markers that incorporate predefined spatial patterns for reliable detection and alignment, RetroLiDAR markers require a mechanism that allows

the detector to identify the start of temporal data within a continuous intensity signal stream. To address this, we design a preamble that is a predefined bit sequence known to both the transmitter and receiver. It is sent before the marker's ID information, and once it is correctly detected, the detector can recognize the beginning of the time-encoded information. Our message format consists of a preamble followed by the encoded message. The marker (using the LCM) first transmits a 7-bit preamble (e.g., "1010101"), followed by the encoded ID information. To comply with the Nyquist sampling theorem, the transmission rate of the LCM is set to half the LiDAR sampling rate. With this packet format, the marker can transmit its information to be detected by the LiDAR.

3.5 Detector Design

Pre-processing LiDAR Signal. At the LiDAR side, we primarily use the temporal intensity map $I(x, y, t)$ (recall §2.1) to detect a RetroLiDAR marker. To ensure proper demodulation in subsequent processing, the temporal distance map of I requires a pre-processing step for abnormal frame elimination. The histogram processing may fail at certain points due to ambient noise or a lack of signal, leading to chaotic signal intensity values that can cause failures. We set the intensity values of all abnormal pixels, $I(x, y, t)$, to $I(x, y, t - 1)$ to mitigate the impact of abnormal frames. This approach maintains the integrity of the temporal distance map and effectively filters out weak noise points (where the majority of statuses are abnormal) by stabilizing the corresponding I values over time. Additionally, we calculate the DC offset $d(x, y)$ for each pixel, which represents the average intensity when the marker is not transmitting.

Preamble Detection. Our preamble detection has two primary goals: (1) *Identify active pixels with marker transmission.* Given the spatial resolution of the LiDAR, not every pixel in its FoV will have an active marker transmitting. We should extract all pixels where a marker is actively transmitting (there may be multiple such pixels, which should not interfere with each other) while marking all other pixels as invalid to prevent false alarms. (2) *Determine the precise timing of the start of transmission.* Since a marker may not begin transmitting immediately due to the duty-cycling design intended to reduce power consumption (§3.3), we need to accurately identify the moment when each marker actually starts transmitting.

- *Computing coefficient and error from least squares.* The basic idea of our preamble detection algorithm is to use a least squares approach to find the position where the received sequence is the most similar to the preamble. Specifically, for each received point $I(x, y, t)$, we evaluate its potential as the preamble start position by first computing the channel distortion coefficient c :

$$c = (\mathbf{p}^T \mathbf{p})^{-1} \mathbf{p}^T \mathbf{s}$$

In this formula, $\mathbf{s} = [I(x, y, t), I(x, y, t + 1), \dots, I(x, y, t + \tau)]^T$ represents the cropped received signal intensity sequence, and \mathbf{p} is a column vector representing the preamble sequence. With the coefficient c , we can calculate the root mean square error between the estimated distorted preamble and the observed signal sequence:

$$e = \sqrt{\sum_{i=1}^{\tau} (s_i - c \cdot p_i)^2}$$

- *Computing likelihood of preamble.* According to the principle of least squares, the absolute value of the coefficient $\frac{|c|}{d}$ indicates the proportion of the modulated signal power in the total power received by the LiDAR, while the ratio $\frac{e}{|c|}$ represents the proportion of error magnitude relative to the signal strength. If a point (x, y) has strong signal power, it is more likely that the pixel is receiving transmissions from a marker because the ambient objects can hardly change its reflectivity so quickly. Similarly, if the error magnitude is small, it is more likely that this point is the start of the preamble sequence. Therefore, we define the likelihood l as:

$$l = \frac{c^2}{e \cdot d}$$

which integrates the metrics and estimates the probability that a given point is the start of the preamble¹.

- *Identifying the start time of transmission.* Suppose the marker is expected to wake up from the duty-cycling design within a duration t_{wake} . We calculate the likelihood l for all (x, y) coordinates and for $0 \leq t \leq t_{\text{wake}}$. We then filter out any pixel (x, y) for which $\forall t, l(x, y, t) < l_{\text{thres}}$, indicating no marker is transmitting in that pixel. For the remaining pixels (i.e., valid pixels), we identify the start of transmission time by sequence correlation:

$$t_{\text{start}}(x, y) = \arg \max_{0 \leq t \leq t_{\text{wake}}} l(x, y, t)$$

corresponding to the peak of $l(t)$.

Demodulation, Decoding, and Ranging. After identifying the preamble start time t_{start} , we demodulate the bits received by each valid pixel starting from $t_{\text{start}} + \tau$ using a binary threshold. Specifically, in our modulation scheme, the bit "0" corresponds to $I(x, y, t) = d(x, y)$, while the bit "1" corresponds to $I(x, y, t) = d(x, y) + c(x, y)$. We set a binary threshold for demodulation, so for each t , the corresponding bit $b(t)$ is determined by:

$$b(t) = \begin{cases} 1, & \text{if } |I(t) - d(x, y)| > |I(t) - (d(x, y) + c(x, y))| \\ 0, & \text{otherwise} \end{cases}$$

Here, $d(x, y)$ is obtained from pre-processing and $c(x, y)$ is the output of preamble detection. If $I(x, y, t)$ is closer to $d(x, y)$, we classify it as "0", and if it is closer to $d(x, y) + c(x, y)$, we classify it as "1". Following bit demodulation, error correction is applied to recover any erroneous bits, provided the error count does not exceed the system's correction capability. Additionally, the marker's distance is estimated by averaging the LiDAR's ranging measurements over the relevant time window:


$$D_{\text{estimated}}(x, y) = \frac{\sum_{t=t_{\text{start}}}^{t_{\text{end}}} D(x, y, t)}{\tau + t_{\text{packet}}}$$

where $t_{\text{end}} = t_{\text{start}} + \tau + t_{\text{packet}}$.

4 LINK BUDGET ANALYSIS

We adopt a link model to analyze the system's performance based on the link model described in PassiveVLC [33], which also utilizes LCM and retroreflector to modulate information. This model expresses the received signal strength at the reader side P_r as a function of several parameters. After omitting uncontrollable or unavailable factors, the simplified relationship reduces to $P_r \propto A_{r_e} / R^4$,

¹In practice, we use $e + 1$ instead of the e to avoid division by zero.



(a) Distance. (b) Size of retroreflector. (c) Integration time. **Figure 5: Experimental validations for link budget analysis.**

where A_{re} is the projected area of the retroreflector and R is the detection distance. To validate the model, we conduct two experiments to quantify the influence of these variables on signal strength. We further investigate the impact of integration time, a LiDAR-specific parameter, on received signal-to-noise ratio (SNR).

Distance. We first evaluate the relationship between retroreflector distance and received signal intensity by varying the marker’s position relative to the LiDAR. As shown in Fig. 5a, for distances beyond 1.5 meters, the signal intensity is inversely proportional to the 3.8th power of the distance, which aligns well with our theoretical model with $R^2 = 0.9826$. At shorter distances, however, minor deviations are observed due to the narrow divergence angle of the retroreflected beam ($\sim 1^\circ$) and the physical separation between the LiDAR’s laser transmitter and the SPAD sensor. Despite this, the signal intensity consistently increases as the distance decreases, confirming the expected trend.

Size of Retroreflector. We further validate the impact of the retroreflector size (A_{re}) on the signal intensity. The results, as shown in Fig. 5b, confirm a strong linear relationship between the retroreflector size and the signal strength, with an R^2 value of 0.9853. The y-intercept observed in the experimental results can be attributed to the influence of the tripod used for the tag in the experiment, which aligns with our expectations. This strong linear correlation supports our model’s predictions and indicates a longer detection distance could be achieved with a larger marker.

Integration Time. The link model aligns well with our LiDAR; however, the LiDAR reader introduces a unique signal reception parameter: integration time. Integration time, the duration during which a SPAD sensor accumulates photon signals from laser pulses (§2.1), directly affects the SNR of LiDAR systems. We conduct experiments to validate the relationship between the LiDAR’s integration time and SNR, as depicted in Fig. 5c. The measured SNR degradation (derived from adjusted sampling intervals with a constant overhead [37] excluded) precisely follows the theoretical $10 \times \log_{10}(k)$ dB loss characteristic of k -times oversampling systems. This quantitative agreement between empirical measurements and communication theory confirms the LiDAR receiver’s compliance with established signal processing principles, thereby validating its potential as a reliable optical communication front-end receiver.

5 IMPLEMENTATION

We implemented a prototype system to validate RetroLiDAR, comprising two major components: the LiDAR and the marker. The LiDAR is based on a VL53L8CX dToF sensor [38] connected to a MCU, running our custom firmware to stream received signals to a PC for advanced signal processing. On the marker side, we designed




Figure 6: RetroLiDAR marker prototype (back and front).

and fabricated a custom frontend board (as shown in Fig. 6), which is connected to an external MCU to handle the logics.

LiDAR Configuration. Our LiDAR setup utilizes the P-NUCLEO-53L8A1 development kit [39], which provides a set of APIs for the VL53L8CX sensor through X-CUBE-TOF1 library [40]. We configure the LiDAR to operate with a stable frame interval of 21 ms², a 4x4 spatial resolution, and in blocking continuous ranging mode. We extend the official firmware by implementing a precise timestamp function using an MCU timer to ensure stable timestamps. The LiDAR sends the frames to an external PC, where further signal processing is performed using Python. As specified in the datasheet [41], a standalone VL53L8CX sensor exhibits a power consumption of 215 mW. Using a Monsoon power monitor [42], we measure the complete development kit’s power consumption at 427 mW during active measurement cycles, including sensor sampling, MCU-based signal processing, and data transmission.

Marker Frontend Circuit. The frontend circuit of our marker is fabricated on a 4.5 cm × 4.5 cm PCB board. The optical receiver of RetroLiDAR marker incorporates a near-infrared sensitive photodiode [43] and an operational amplifier (OPA320) [44] configured as a transimpedance amplifier with a 2 MΩ resistance (providing a first-stage amplification factor of 2,000,000) and a 3 pF feedback capacitor. A signal pattern filter comprising two series RC high-pass filters processes the photodiode output. The filtered signal is then amplified by another OPA320 operational amplifier, rectified using a SMS7630-079LF diode [45], and finally compared against a pre-defined threshold (see §6.5) using a comparator [46]. The comparator’s output is utilized to trigger the MCU from low power detection mode to active transmission mode. The end-to-end detection latency is less than 0.4 ms based on our circuit model simulation on ADS [47]. The liquid crystal driver circuit uses a comparator, which outputs a 5V voltage to the liquid crystal electrodes whenever the GPIO output from the MCU is at a high level.

²Although the documentation indicates support for up to 60 Hz, we could not achieve this frame rate due to I2C bus bandwidth limitations and API overhead.




Figure 7: Experimental setups.

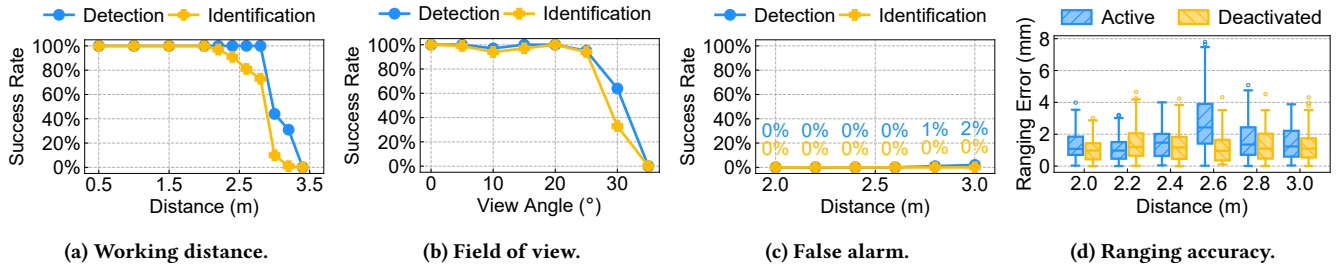


Figure 8: Marker reading performances.

Table 2: Power consumption under different modes.

Mode	Sleep	Detection	Active
Power (mA)	0.001	3.98	12.7
Duty cycle	99%	1%	Task dependence (e.g., < 0.1%)

Power Consumption. We fabricate a fully functional RetroLiDAR marker to conduct end-to-end power consumption analysis. The power consumption measurements are shown in Tab. 2. In detection mode, the marker consumes 3.98 mA (2 mA frontend + 1.98 mA MCU) during quiescent operation. Upon detecting a LiDAR signal, the MCU is fully activated, resulting in an active current of 12.7 mA. Based on a 90 mAh coin cell battery [48], a detection duty cycle of 1 ms / 100 ms, and a query frequency of once per hour, the marker can operate for two months. The power consumption could be further optimized to approximately 100 μ A with ultra-low-power operational amplifier (e.g., MAX9917 [49]) and low-power MCU, extending the battery life to over two years.

Cost Analysis. The LiDAR development kit used in our system cost \$36.25 for each, while a discrete VL53L8CX dToF sensor is priced at \$9.09. For the marker, the electronic components contribute \$9.91 to the Bill of Materials (BOM) cost, with an additional \$112 allocated for optical components. The optical component breakdown includes \$100 for a customized LCM in the near-infrared band, \$9 for the photodiode, and \$3 for the infrared spectrum filter. Our LCM is a custom non-commercial component, whereas the PD and spectrum filter are commercially available infrared devices. We believe that once infrared LCMs become commercially available, their price will be comparable to the other two components (under \$10), reducing the total marker cost to under \$40.

6 EVALUATION

6.1 Experimental Setup

We evaluate RetroLiDAR using our prototype implementation with a testbed shown in Fig. 7. The LiDAR and marker are mounted on tripods to conduct distance experiments. We also utilize two optical benches for experiments that require precise angle control. Unless otherwise stated, the marker (3×3 cm²) is positioned directly in the center of the LiDAR’s field of view at a distance of 1 meter, with no yaw misalignment, and in typical office environments where illumination levels measured around 200 lux at night with lights on and increased to \sim 500 lux during daytime due to natural light ingress through adjacent windows. Our primary metrics for evaluation are two metrics:

- **Detection rate.** The detection rate quantifies the ability of the LiDAR to detect the marker’s presence by successfully decoding its preamble sequence.
- **Identification rate.** The identification rate measures the LiDAR’s capability to successfully decode the complete packet and correctly obtain the marker’s ID.


We define normal marker operation as achieving an identification rate exceeding 80%, because this ensures a >99% probability of correct identification within three consecutive reads. For each experimental condition, a marker transmits 10 random packets (1.3 seconds per packet) following our packet format, repeating this process 10 times to ensure statistical significance.

6.2 Marker Reading

Working Range. We first conduct detection and identification rate tests under varying distances and angles. As illustrated in Fig. 8a, RetroLiDAR operates effectively at distances of up to 2.6 meters and consistently achieves a 100% success rate in identifying the marker’s ID within 2.0 meters. Detection performance remains at 100% up to a maximum distance of 2.8 meters. We then proceed to test the angular FoV for RetroLiDAR. We place the marker at different positions within the LiDAR’s FoV with a fixed distance of 1 meter, recording the relationship between the view angle (the angle between the marker’s position and the center of the LiDAR’s field of view) and the success rates. As shown in Fig. 8b, RetroLiDAR demonstrates effective detection and identification performance within a working angle of ± 25 degrees, allowing a sufficient FoV.

False Alarm. When no active marker is present within its field of view, the LiDAR still output distance and intensity maps, which could potentially lead to false alarms. To assess the false alarm rates, we conduct an experiment where the LiDAR attempts to read a marker that is intentionally deactivated and does not modulate signals. Fig. 8c shows that RetroLiDAR exhibited a false detection rate of only 0.5% across different distances due to ambient noise. More importantly, no false identification is observed during the experiment. These results demonstrate the effectiveness of our preamble detection algorithm. Furthermore, due to the error detection capability of the coding scheme, the likelihood of decoding a valid code after a false detection is negligible.

Ranging Accuracy. Given the widespread acceptance of LiDAR ranging as ground truth, we evaluate RetroLiDAR’s ranging accuracy through statistical analysis of repeated distance measurements. We compare the standard deviation of estimated distances during active marker detection (LCM modulation enabled) against silent




(a) Distance ($3 \times 3 \text{ cm}^2$). (b) Distance ($1 \times 1 \text{ cm}^2$).
Figure 9: Performance comparison with AprilTag.

marker baselines (LCM deactivated) across varying operational distances. As shown Fig. 8d, in all cases except at a distance of 2.6 meters, the median ranging error remains below 2 mm, confirming that our design fully leverages LiDAR’s inherent ranging precision regardless of marker activity. At 2.6 meters with the marker active, we observe a slight increase in error due to nearby human movement, yet the median error remains under 2.5 mm. Additionally, we test LiDAR performance on a bare PCB of the same size (without optical components like retroreflectors or LCM) to represent its typical ranging capabilities on generic surfaces. At a distance of 2.0 meters, the median ranging error is 1.0 mm with a maximum error of 4.2 mm, closely matching the results achieved with our marker. These results confirm that RetroLiDAR not only enables marker identification but also maintains LiDAR’s high ranging accuracy, in contrast to the results in Tab. 1.

Table 3: Identification rates of individual markers when simultaneously read.

Marker arrangement	1	2	3	4	5
Marker #1	97%	96%	95%	86%	92%
Marker #2	98%	93%	85%	90%	98%
Marker #3	91%	92%	98%	87%	90%

Simultaneous Reading of Multiple Markers. The SPAD LiDAR of RetroLiDAR has spatial resolution capabilities. It is able to simultaneously read information from multiple pixels in space, thereby enabling the simultaneous reading of multiple markers. To validate this capability, we place three markers at varying distances (1.0 to 2.0 meters) and angles, arranging them in five different spatial configurations for the LiDAR to read. Experimental results demonstrate consistent 100% detection rates across all arrangements, with individual marker identification rates exceeding 85% (see Tab. 3). Since signals from spatially separated markers do not interfere with each other, the system does not need to adjust the emitted light intensity of the LiDAR and inherently avoids near-far effects. However, identification rates varies among individual markers, reflecting a performance decrease compared to single-marker scenarios. We attribute the performance drop to the current prototype’s limited 4×4 spatial resolution, which increases susceptibility to interference. Future implementations leveraging higher-resolution LiDAR systems are expected to address this limitation.



(c) Ranging error ($3 \times 3 \text{ cm}^2$). (d) Identification Rate vs. Distance (m).
Figure 10: Distance comparison with $3 \times 3 \text{ cm}^2$ QR Code.

6.3 Comparison with Visual Markers

We design an experiment to evaluate the identification and ranging performance of two mainstream visual fiducial markers, AprilTag and QR Code, in comparison with RetroLiDAR. We use three smartphones, spanning the high-end, mid-range, and entry level models with their detailed specifications listed in Tab. 4, to scan the visual markers. The smartphones are calibrated using the OpenCV chessboard calibration method [50]. We configure the markers to encode a number of bits as close as possible to the same amount in each marker. The markers are placed at varying distances and angles relative to the smartphone main camera³. The cameras record video at FHD resolution and 30 FPS, which is the default setting for video capturing APIs. For each setup, we calculate the average identification rate and ranging error based on 1000 random video frames using corresponding detectors.

Table 4: Parameters of visual marker systems.

Smartphone	Model	Release year	Launch price
High-end	iPhone 15 Pro Max	2023	\$1199
Mid-range	Samsung Galaxy S21	2021	\$699
Entry-level	OnePlus 7T	2019	\$599
Marker	Tag family	Dictionary size	Detector
AprilTag	tagStandard41h12	2115	Apriltags 3 [51]
QR Code	Version 1	16 bytes	zbar [52]

AprilTag Identification. As shown in Fig. 9a, for a $3 \times 3 \text{ cm}^2$ AprilTag, the maximum reliable identification range is 2.2 meters on entry-level smartphones and extends to 2.4 meters on high-end models. However, when the tag size is reduced to 1 cm^2 , the identification range for standard smartphone cameras significantly decreases to less than 0.9 meters (Fig. 9b). In contrast, RetroLiDAR achieves a reliable identification range of 1.3 meters with a tiny 1 cm^2 marker, representing a 44% improvement (the details are to be illustrated in §6.4). This result demonstrates RetroLiDAR’s superior performance especially with smaller markers. Notably, these results are obtained under typical indoor lighting, where the performance of visual markers degrades in dark environments. Meanwhile, RetroLiDAR maintains reliable (and even better) identification in complete darkness, as will be further elaborated in §6.4.

³All the main cameras have a FoV similar to that of our LiDAR (± 25 degrees). While the reading distance can be extended using telephoto lenses, this approach reduces the FoV. To ensure a fair comparison, our experiments maintain the same FoV for both cameras and LiDAR.

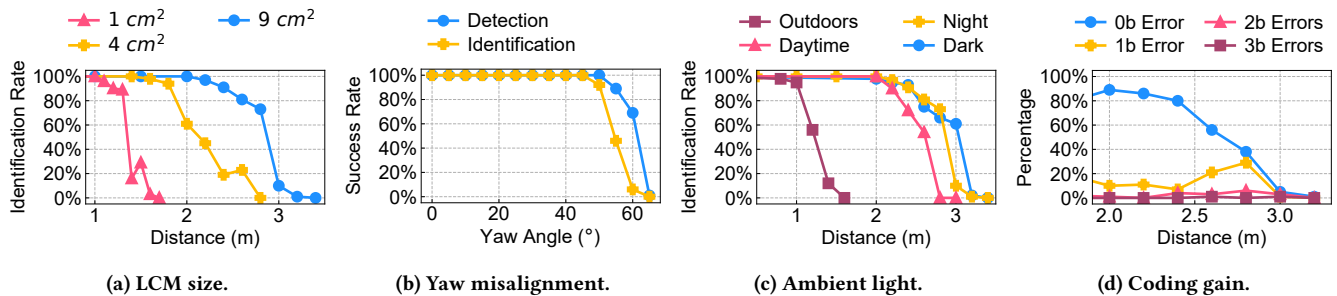


Figure 11: Microbenchmarks.

AprilTag Ranging. Regarding ranging accuracy, Fig. 9c highlights the limitations of visual fiducial markers, which are constrained by the sensing capabilities of standard cameras. At a favorable distance of 1.0 meter, the median ranging error for a $3 \times 3 \text{ cm}^2$ AprilTag is under 4 mm, consistent with literature reports. However, the errors can still exceed 10 mm in certain cases. As the distance increases, the ranging accuracy deteriorates significantly. For example, even with high-end smartphones, at a distance of 2.4 meters (the maximum reliable identification range), the median error increases to 11.7 mm, with the maximum error reaching 156.2 mm. These findings align with prior benchmarks [16]. In contrast, RetroLiDAR consistently achieves sub-centimeter-level ranging accuracy at various distances (Fig. 8d), offering a clear improvement over visual markers.

QR Code. As shown in Fig. 10, QR Codes with a size of $3 \times 3 \text{ cm}^2$ exhibit significant limitations in identification range, with reliable detection confined to less than 1 meter. Even at closer distances (around 0.5 m), the identification rate fails to approach 100% due to distortion at the edges of the camera’s FoV. In contrast, RetroLiDAR extends this distance to 2.6 meters, achieving a substantial 2.6x improvement. This result highlights the challenges faced by visual markers in maintaining identification performance with limited marker sizes. In terms of ranging accuracy, the performance degradation aligns closely with those observed for AprilTags. Detection failures and ranging faults become frequent beyond 1 meter, rendering QR Codes unable to provide any reliable ranging data at greater distances. In contrast, RetroLiDAR maintains high identification reliability and sub-centimeter ranging precision.

6.4 Microbenchmarks

LCM Size. We investigate the impact of LCM size on identification performance by progressively covering portions of the marker’s LCM. As shown in Fig. 11a, although the reliable identification distance decreases as the area is reduced, the relationship is nonlinear. Even with a tiny 1 cm^2 marker, RetroLiDAR can reliably identify it at a distance of 1.3 meter. This result aligns with our link budget analysis (§4) and highlight RetroLiDAR’s potential for applications requiring extremely small markers.

Angular Misalignment. When the marker is placed not directly facing the LiDAR (*i.e.*, experiencing yaw angular misalignment), RetroLiDAR’s performance may degrade due to the reduced projected area facing the LiDAR and potential impacts on the retro-reflector’s directional properties. We design an experiment to assess the effect of yaw misalignment. Fig. 11b reveals that the our marker

can tolerate approximately ± 50 degrees of yaw angular misalignment before a rapid decline in signal reflection capabilities, with complete failure occurring beyond 60 degrees. We believe that the working range already covers the majority of fiducial marker applications, where markers are typically installed in fixed orientations. A better physical marker design that explores multi-faceted marker geometries may further broaden the effective angular range for scenarios requiring extreme angular flexibility.

Ambient Light Condition. To evaluate RetroLiDAR’s performance under various ambient light conditions, we conduct experiments under three indoor scenarios: daytime with sunlight entering through windows, nighttime with artificial lighting, and complete darkness. Fig. 11c shows that even during the daytime, the reduction in working distance compared to darkness is only about 20 cm. This result demonstrates RetroLiDAR’s resilience to varying indoor illumination. Additionally, although RetroLiDAR is primarily designed for indoor scenarios, we also test its performance outdoors during the daytime. Although the sensor fails under direct sunlight, it can still achieve reliable readings at distances over 1 meter when not in direct sunlight. The outdoors result further highlights RetroLiDAR’s capability to withstand varying ambient light conditions.

Error Correction Coding Gain. RetroLiDAR incorporates error correction coding to ensure robust performance. To evaluate its importance, we analyze the raw demodulation results from correctly decoded readings and compare them with the transmitted sequences, quantifying the percentage of readings exhibiting varying bit error counts. As shown in Fig. 11d, even when the identification rate is 100%, only 89% of the readings are correct without the need for error correction. As the distance increases, the percentage of readings with bit errors continued to rise; at 2.8 meters, 48% of the correct readings contain bit errors. More detection errors occurred at greater distances. Our coding scheme effectively helps increase the correct identification rate to nearly 100%, demonstrating its crucial role in maintaining high identification rate.

6.5 Wake-up Circuit Subsystem Performance

A RetroLiDAR marker is equipped with a wake-up circuit that enables targeted responses when queried by LiDAR. We first measure the strength of the LiDAR signal received by the PD at various distances. Following this, we validate the effectiveness of our filter design. By combining these measurements and optimally tuning the threshold, we achieve reliable and sensitive wake-up triggering.

Received Signal Energy. We first measure the signal strength that our circuit could receive from the LiDAR at various distances

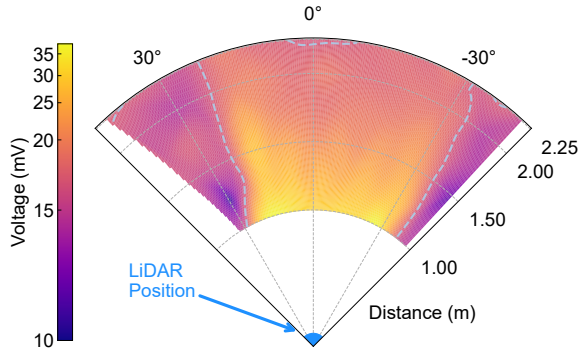


Figure 12: Signal intensity of LiDAR emissions received by the marker at different positions.

and angles. Specifically, we measured the voltage root mean square (VRMS) output of the PD while varying the distance and orientation of the LiDAR. Our experimental results (Fig. 12) indicate that within a distance of 2.25 meters and an angle of 55°, the marker circuit can receive energy from the LiDAR that exceeds the circuit’s noise level. Beyond this range, the signal strength falls below the noise level, making it impossible for the circuit to function correctly. Notably, the relationship between signal strength and distance follows an inverse square law, consistent with the predictions from our link budget analysis model (§4).

Table 5: Voltage output of PD under different conditions with and without filters.

Voltage (mV)	Dark	Indoors	Outdoors	W/ LED	W/ LiDAR
No filter	632.72	5132.3	5182.3	5163.4	5072.0
Spectrum filter	64.232	191.82	5067.2	594.33	653.42
Signal pattern filter	50.111	50.737	38.149	50.201	50.239
Both filters	14.812	16.443	36.695	16.895	32.413

Effectiveness of Filters. Our circuit incorporates a dual-filter design to isolate the LiDAR signal from ambient noise. We evaluate the performance of both filters by measuring the VRMS output from the PD under different light conditions, both with and without filters. The results are summarized in Tab. 5. A spectrum filter is used to prevent saturation; without it, the PD would saturate even under typical indoor lighting, rendering the reception of any signal impossible. After passing through the spectrum filter, a signal pattern filter is employed to specifically extract the high-frequency pulses sent by the LiDAR. Without the signal pattern filter, the output would be sensitive to other infrared light sources (e.g., the LED of an infrared camera, denoted as W/ LED in Tab. 5). After applying both filters, the output signal is selectively tuned to the infrared frequency and the transmission pattern of the LiDAR. These results confirm that the filters effectively extract the characteristic signals of the LiDAR, ensuring reliable operation of the wake-up circuit.

Detection Distance. The measurements of the PD output can be used to adjust the threshold for wake-up triggering. We set the triggering threshold to the noise level under typical indoor lighting conditions, with both filters applied, plus an additional 1

mV guard value. By combining the threshold with the previously tested PD output levels under normal LiDAR illumination at various distances, we determine that our design could achieve wake-up triggering at a distance of 2.25 meters. Given that this distance might be shorter than the effective marker reading range, we conduct an additional experiment using the VL53L4CX LiDAR model [53], which has the same transmission power but a narrower FoV of 18°. Our experiments verified that this LiDAR could successfully trigger wake-up at distances up to 3.75 meters, further validating the effectiveness of our wake-up circuit design.

7 USE CASE STUDY

We demonstrate RetroLiDAR’s potential in embodied AI applications through two proof-of-concept use cases, highlighting its ability to enhance perception fidelity in real-world scenarios.

7.1 Coffee-making Tasks

In embodied AI applications requiring physical interaction, precise object identification coupled with spatial relationship comprehension forms the foundation of reliable task execution. Consider a coffee preparation scenario (Fig. 13a) where an agent must simultaneously detect the coffee beans, machine, and mug, as well as understand their spatial relationships. By deploying RetroLiDAR markers on these objects, we enable millimeter-accurate identification and spatial reasoning through LiDAR’s inherent depth perception. We would like to mention that accuracy is crucial for the overall execution efficiency. For example, common robotic arms (e.g., MyCobot [54]) typically need 2-5 seconds for each movement. Therefore, inaccurate measurements of vision-based positioning (resulting in iterative adjustments) critically constrain the performance of task execution. To quantify this enhancement, we conduct 30 trials with varying LiDAR positions/orientations, measuring the duration to detect and localize all three markers. The results show a mean multi-object identification and ranging latency of 5.3 seconds, including 3 seconds for one-shot robotic arm movement. In contrast, an AprilTag-based system requires position recalibration when localization errors exceed 1 cm (the possibility is 36% in the experiments in Fig. 9c). With three markers, the inaccurate localization leads to an average of 1.7 additional measurement cycles per task due to recalibration, extending the overall task finish time to 8.1 seconds (2.7 measurements \times 3 seconds per movement). This inefficiency becomes more severe when the accuracy requirement is higher. In summary, RetroLiDAR achieves one-shot localization with millimeter-level accuracy, eliminating the need for iterative recalibration and reducing task execution latency.

7.2 Drone Operation

Precise real-time localization is critical for embodied agents to reliably operate machines and interact with their environments. For example, drones can operate in confined indoor environments such as warehouses, factories, and office buildings, where they typically fly at relatively low speeds but require continuous, high-accuracy positional updates to avoid collisions with static structures (e.g., walls, ceilings) and dynamic obstacles (e.g., other drones). In this use case, RetroLiDAR leverages its inherent advantages in ranging precision and environmental robustness to provide superior localization data. To validate this capability, we conduct an experiment using a DJI Air 2S drone equipped with a marker beneath its rotor,

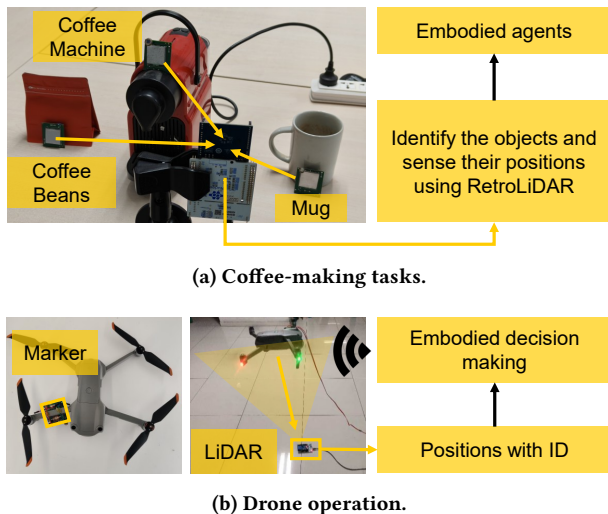


Figure 13: Use case study.

hovering approximately 1 meter above a ground-mounted LiDAR sensor (see Fig. 13b). Despite turbulence-induced positional drift (>20 cm displacement) and intermittent self-occlusion caused by rotor movements, RetroLiDAR achieves a 91% heartbeat signal detection rate within 10-second intervals over a continuous 260-second trial, maintaining sub-1 cm ranging accuracy. In contrast, visual fiducial markers on drones exhibit significantly higher localization errors (e.g., >0.16 meters [55]). This performance gap highlights RetroLiDAR’s unique suitability for dynamic environments where vision-based approaches struggle to maintain fidelity.

8 DISCUSSION

Cost for Large-scale Deployment. Our marker is more expensive than conventional printed fiducial markers, and we acknowledge that this represents a significant limitation. However, we believe that for embodied AI applications in the current stage, a cost of tens of dollars is not the primary concern. Instead, the high fidelity offered by our system may be more attractive. On the other hand, mobile LiDAR is currently integrated into high-end smartphones, but we anticipate that with advancements in semiconductor manufacturing and the reduction of R&D costs, the price of LiDAR will decrease, making it accessible for mainstream robots.

LiDAR Capabilities. Mobile LiDAR technology is still in its early stages, but it has made significant strides in recent years. However, broader applications still require improvements in frame rate and pixel resolution, which are critical for faster data acquisition and enhanced mobility support. Sensor fusion strategies that integrate LiDAR with complementary modalities (e.g., RGB cameras or high-frame-rate neuromorphic cameras) could also enhance the spatial/temporal resolution beyond LiDAR’s inherent limitations. Advancements in spatial resolution could further enable 3D pose estimation through multi-pixel marker designs, where multiple positional parameters of a marker could be simultaneously extracted and angular information could be calculated. Furthermore, increasing the detection distance against ambient noise would enable more effective use in outdoor daylight environments, which still remains a common limitation of LiDAR systems. We would like to mention

that better APIs are equally essential. For instance, while we experimented with Apple iPhone’s LiDAR that is considerably more powerful, the provided API lacked any necessary information and flexibility for our use case. Finally, future LiDAR systems that support active control of emitted signals could enable more complex modulation schemes, allowing markers to filter out noise and select wake-up signals more effectively.

Better Mobility Support. RetroLiDAR provides partial mobility support, functioning when the marker experiences slight movements, as demonstrated in our use case study. For embodied AI tasks, a reading time of less than 1 second is unlikely to be the primary bottleneck in overall task completion. The key limitation preventing further enhancement of mobility support is the insufficient frame rate, which leads to blurred reception of LiDAR signals during movement. Increasing the frame rate, which is currently constrained by the capabilities of commercial devices, would significantly mitigate the impact of mobility on system performance.

9 RELATED WORK

Visual Fiducial Markers. Visual fiducial markers are artificial landmarks designed to serve two primary functions using cameras: identification (~ 10 bits information) and position estimation. Various types of markers have been proposed, such as ARTag [10], AprilTag [11], ArUco [12], ChromaTag [63], and STag [64]. They have a broad range of applications including augmented reality [65], robot navigation [66, 67], drone landing [68, 69], large-scale 3D printing [70], surgery [71], and animal behavior tracking [72]. Beyond the conventional marker systems that utilize RGB cameras to read printed markers, recent research has explored more advanced hardware designs to achieve additional functionalities such as single-pixel imaging for privacy-preserving reading [73], leveraging the optical image stabilization module of a smartphone camera to reconstruct 3D images [74], utilizing the dual cameras on mobile devices for depth estimation [75], designing Moiré patterns for marker pose measurement [76, 77], using birefringent nature of retroreflective tags for accurate angular measurement [78, 79], and designing 3D markers [80] or LED beacons [81] for underwater navigation. RetroLiDAR extends the capabilities of fiducial markers by utilizing liquid crystals to embed information in the time domain and employing LiDAR for reading.

Wireless Marker Systems. Wireless marker systems leverage various wireless communication technologies to enable identification and localization of tags. RFID currently stands as the most ubiquitous wireless marker technology in commercial deployments. Numerous systems have been developed for RFID-based localization [82, 83] and human-computer interaction [84]. However, due to limitations in spatial and angular resolution, RFID localization systems often require extensive ranges of movement of the reader [84], large antenna apertures [85], or costly devices [61], rendering them unsuitable for practical use. Bluetooth [60] and Ultra-Wideband markers [86] are also widely used for indoor identification and localization, but their localization errors are 10-100x higher than vision-based markers. Recently, backscatter technologies such as Wi-Fi [87], LTE [88], LoRa [89], and mmWave [90, 91] backscatter have gained popularity. Despite their rapid development, these technologies are still in the early stages [92] and face challenges in

Table 6: Comparison of different marker systems.

System	Technology	Working Range	Localization Error (Median)	Reader Price ^a	Marker Price ^a	Marker Size	Environmental Robustness ^b
RetroLiDAR (Ours)	LiDAR + LCM	2.6 m	1 mm	\$36	< \$40	45 × 45 mm²	Medium
LiDARTag [25]	LiDAR	14 m	2.8 mm	\$7500	~ \$10	1.2 × 1.2 m ²	Medium
IILFM [26]	LiDAR	4 m	0.11 m	\$600	~ \$0.1	0.17 × 0.17 m ²	Medium
L-PR [28]	LiDAR	10 m	N/A	\$600	~ \$1	0.69 × 0.69 m ²	Medium
AprilTag [11]	Camera	2.4 m	11.7 mm	-	~ \$0.05	30 × 30 mm ²	Low
QR Code	Camera	1.0 m	8.1 mm	-	~ \$0.05	30 × 30 mm ²	Low
SuperSight [56]	mmWave	8 m	14.8 mm	\$3600	\$32	40 × 40 mm ²	High
Hawkeye [57]	mmWave	25 m	2.08 m	\$200	~ \$30	66 × 66 mm ²	High
Apple AirTag [58]	UWB ^c	10 m	0.26 m	-	\$29	31.9 × 31.9 mm ²	High
ULoc [59]	UWB	8.5 m	0.12 m	~ \$150	~ \$30	37 × 16 mm ²	High
Bloc [60]	Bluetooth	6 m	0.86 m	N/A	N/A	N/A	High
RF-Chord [61]	UHF RFID	6 m	0.14 m	~ \$7000	~ \$0.1	80 × 80 mm ²	High
MetaSight [62]	UHF RFID	7 m	0.15 m	~ \$1900	~ \$0.1	48 × 13 mm ²	High

^a For prices not explicitly reported in the literature, we provide estimates based on our analysis and comparable technologies.

^b We define environmental robustness through three tiers: (1) High - maintaining full functionality under illumination variations and adverse weather conditions (fog, rain, *etc.*), (2) Medium - sustaining reliable operation across a wide range of illumination levels, and (3) Low - requiring strictly controlled and consistent lighting conditions to ensure proper functionality.

^c AirTag utilizes Bluetooth for coarse-grained localization in addition to UWB.

being applied to applications that require high fidelity. Tab. 6 provides a systematic comparison between RetroLiDAR and existing marker solutions using different media.

Liquid-crystal Wireless Systems. Liquid-crystal wireless systems leverage liquid crystals to modulate visible light for data communication, primarily through backscatter techniques. Various modulation schemes, including time-domain OOK [33, 93], pulse amplitude modulation [94], time-polarization collaborative modulation [95], and multi-user MIMO in the pulse feature domain [96], have been proposed to utilize the visible light backscatter channel using LCM and retroreflector. Beyond backscatter systems, there are also transmissive systems that do not rely on retroreflectors but modulate transmissive ambient light (*e.g.*, sunlight) using frequency-domain modulation [97] and color-based modulation [98, 99]. These technologies have enabled diverse applications, including indoor localization [100–102], human-robot collaboration [103], local information provision [104], joint illumination and communication [105], and vehicle-to-everything (V2X) networks [106–108], expanding the scope of liquid crystal wireless systems in both research and practical implementations.

Mobile LiDAR Applications. With the commercialization of SPAD LiDAR systems in mobile devices, an increasing number of researchers are leveraging this emerging ToF sensor for various research endeavors. The precise ranging capabilities of LiDARs enable applications such as 3D model reconstruction [109], geographic information capture [110, 111], indoor spectral mapping [112], and mmWave coverage optimization [113]. Their small form factor and low power consumption make mobile LiDARs particularly suitable for integration into small robots, facilitating pose estimation [114] and robot detection [115]. In addition to spatial perception, mobile LiDAR systems also enables innovative applications including material classification [116, 117], recovery of planar geometry and albedo [118, 119], and liquid testing based on laser

speckle phenomena [120]. These research efforts have expanded the sensing capabilities of mobile devices.

10 CONCLUSION

In this paper, we have demonstrated a capability where computer vision (CV) falls short – precise ranging – while wireless techniques can be leveraged to achieve superior performance. Despite CV’s remarkable progress in general tasks, it remains fundamentally limited in ranging capabilities due to the inherent constraints of the sensor itself. By using LCMs to encode information through dynamic temporal variations rather than static geometric features, our system achieves extended reading distances (2.6× improvement) and higher ranging accuracy (85% error reduction) with a tiny $1 \times 1 \text{ cm}^2$ marker. This principle parallels RF backscatter’s effectiveness through intentional modulation, contrasting with chipless RFID’s dependence on bulky antenna patterns for spatial differentiation. Such miniaturization enables centimeter-scale markers to operate unobtrusively within environments while maintaining high detection and ranging fidelity, paving the way for advanced sensory spatial data collection in future applications.

ACKNOWLEDGMENTS

We sincerely thank the anonymous shepherd and reviewers for their insightful critique and constructive feedback, which have significantly improved the quality of this paper. We are also grateful to Zhe Ou for his contributions to the early versions of this work, Vigneshwaran Swaminathan and Abhishek Kumar Srivastava for their valuable comments on infrared LCM design, and Yutong Liu for his assistance with visual marker data collection. This work is supported by National Key Research and Development Plan, China (Grant No.2023YFB2903902) and National Natural Science Foundation of China (Grant No.62272010 and 62061146001). Chenren Xu (chenren@pku.edu.cn) is the corresponding author.

REFERENCES

- [1] Zane Durante, Qiuyuan Huang, Naoki Wake, et al. Agent ai: Surveying the horizons of multimodal interaction. *arXiv:2401.03568*, 2024.
- [2] Laura Londoño, Juana Valeria Hurtado, Nora Hertz, Philipp Kellmeyer, Silja Voeneky, and Abhinav Valada. Fairness and bias in robot learning. *Proceedings of the IEEE*, 2024.
- [3] Yang Liu, Weixing Chen, Yongjie Bai, Xiaodan Liang, Guanbin Li, Wen Gao, and Liang Lin. Aligning cyber space with physical world: A comprehensive survey on embodied ai. *arXiv:2407.06886*, 2024.
- [4] Yang Bai, Nakul Garg, and Nirupam Roy. Spidr: Ultra-low-power acoustic spatial sensing for micro-robot navigation. In *ACM MobiSys*, 2022.
- [5] Jiafei Duan, Samson Yu, Hui Li Tan, Hongyuan Zhu, and Cheston Tan. A survey of embodied ai: From simulators to research tasks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2022.
- [6] Yang Liu, Xinsuai Song, Kaixuan Jiang, Weixing Chen, Jingzhou Luo, Guanbin Li, and Liang Lin. Multimodal embodied interactive agent for cafe scene. *arXiv:2402.00290*, 2024.
- [7] Georges Chahine and Cédric Pradalier. Survey of monocular slam algorithms in natural environments. In *IEEE Conference on Computer and Robot Vision*, 2018.
- [8] Bharat Joshi, Sharmin Rahman, Michail Kalaitzakis, et al. Experimental comparison of open source visual-inertial-based state estimation algorithms in the underwater domain. In *IEEE/RJ IROS*, 2019.
- [9] Adam Polevoy, Craig Knuth, Katie M Popek, and Kapil D Katyal. Complex terrain navigation via model error prediction. In *IEEE ICRA*, 2022.
- [10] Mark Fiala. Arttag, a fiducial marker system using digital techniques. In *IEEE CVPR*, 2005.
- [11] Edwin Olson. Apriltag: A robust and flexible visual fiducial system. In *IEEE ICRA*, 2011.
- [12] Sergio Garrido-Jurado, Rafael Muñoz-Salinas, Francisco José Madrid-Cuevas, and Manuel Jesús Marín-Jiménez. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition*, 2014.
- [13] Jingpei Lu, Florian Richter, and Michael C Yip. Pose estimation for robot manipulators via keypoint optimization and sim-to-real transfer. *IEEE RA-L*, 2022.
- [14] Ali Tourani, Hriday Bavle, Deniz İşınur Avşar, Jose Luis Sanchez-Lopez, Rafael Muñoz-Salinas, and Holger Voos. Vision-based situational graphs exploiting fiducial markers for the integration of semantic entities. *Robotics*, 2024.
- [15] Christopher E Mower, Yuhui Wan, Hongzhan Yu, Antoine Grosnit, Jonas Gonzalez-Billandon, Matthieu Zimmer, Jinlong Wang, Xinyu Zhang, Yao Zhao, Anbang Zhai, et al. Ros-llm: A ros framework for embodied ai with task feedback and structured reasoning. *arXiv:2406.19741*, 2024.
- [16] Michail Kalaitzakis, Brennan Cain, Sabrina Carroll, Anand Ambrosi, Camden Whitehead, and Nikolaos Vitzilaios. Fiducial markers for pose estimation: Overview, applications and experimental comparison of the arttag, apriltag, aruco and stag markers. *Journal of Intelligent & Robotic Systems*, 2021.
- [17] Capturing depth using the LiDAR camera. https://developer.apple.com/documentation/avfoundation/additional_data_capture/capturing_depth_using_the_lidar_camera.
- [18] Apple Vision Pro. <https://www.apple.com/apple-vision-pro/>.
- [19] Weijing Shi and Raj Rajkumar. Point-gnn: Graph neural network for 3d object detection in a point cloud. In *IEEE/CVF CVPR*, 2020.
- [20] Chenhang He, Hui Zeng, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Structure aware single-stage 3d object detection from point cloud. In *IEEE/CVF CVPR*, pages 11873–11882, 2020.
- [21] Lue Fan, Ziqi Pang, Tianyuan Zhang, Yu-Xiong Wang, Hang Zhao, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Embracing single stride 3d object detector with sparse transformer. In *IEEE/CVF CVPR*, 2022.
- [22] Xin Lai, Yukang Chen, Fanbin Lu, Jianhui Liu, and Jiaya Jia. Spherical transformer for lidar-based 3d recognition. In *IEEE/CVF CVPR*, 2023.
- [23] Jiachen Sun, Qingzhao Zhang, Bhavya Kailkhura, Zhiding Yu, Chaowei Xiao, and Z Morley Mao. Benchmarking robustness of 3d point cloud recognition against common corruptions. *arXiv preprint:2201.12296*, 2022.
- [24] Yinpeng Dong, Caixin Kang, Jinlai Zhang, Zijian Zhu, Yikai Wang, Xiao Yang, Hang Su, Xingxing Wei, and Jun Zhu. Benchmarking robustness of 3d object detection to common corruptions. In *IEEE/CVF CVPR*, 2023.
- [25] Jiunn-Kai Huang, Shoutian Wang, Maani Ghaffari, and Jessy W Grizzle. Lidartag: A real-time fiducial tag system for point clouds. *IEEE RA-L*, 2021.
- [26] Yibo Liu, Hunter Schofield, and Jinjun Shan. Intensity image-based lidar fiducial marker system. *IEEE RA-L*, 2022.
- [27] Yusen Xie, Lei Deng, Ting Sun, Yeyu Fu, Jian Li, Xinglong Cui, Hanxi Yin, Shuixin Deng, Junwei Xiao, and Baohua Chen. A4lidartag: Depth-based fiducial marker for extrinsic calibration of solid-state lidar and camera. *IEEE RA-L*, 2022.
- [28] Yibo Liu, Jinjun Shan, Amaldev Haridevan, Shuo Zhang, and Kejian Lin. L-pr: Exploiting lidar fiducial marker for unordered low overlap multiview point cloud registration. *arXiv:2406.03298*, 2024.
- [29] Tingting Liang, Hongwei Xie, Kaicheng Yu, Zhongyu Xia, Zhiwei Lin, Yongtao Wang, Tao Tang, Bing Wang, and Zhi Tang. Bevfusion: A simple and robust lidar-camera fusion framework. *NeurIPS*, 2022.
- [30] Yingwei Li, Adams Wei Yu, Tianjian Meng, et al. Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. In *IEEE/CVF CVPR*, 2022.
- [31] Xin Li, Tao Ma, Yuenan Hou, Botian Shi, et al. Logonet: Towards accurate 3d object detection with local-to-global cross-modal fusion. In *IEEE/CVF CVPR*, 2023.
- [32] Rafael Marques Claro, Diogo Brandão Silva, and Andry Maykol Pinto. Artuga: A novel multimodal fiducial marker for aerial robotics. *Robotics and Autonomous Systems*, 2023.
- [33] Xieyang Xu, Yang Shen, Junrui Yang, Chenren Xu, Guobin Shen, Guojun Chen, and Yunzhe Ni. PassiveVLC: Enabling practical visible light backscatter communication for battery-free IoT applications. In *ACM MobiCom*, 2017.
- [34] Alessandra Spreafico, Filiberto Chiabrando, L Teppati Losè, and F Giulio Tonolo. The ipad pro built-in lidar sensor: 3d rapid mapping tests and quality assessment. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2021.
- [35] F Fischer, Luca Romoli, and R Kling. Laser-based repair of carbon fiber reinforced plastics. *CIRP annals*, 2010.
- [36] Binary Golay code. https://en.wikipedia.org/wiki/Binary_Golay_code.
- [37] A guide for using the VL53L8CX, low-power, high-performance Time-of-Flight multizone ranging sensor. https://www.st.com/resource/en/user_manual/um3109-a-guide-for-using-the-vl53l8cx-lowpower-highperformance-timeofflight-multizone-ranging-sensor-stmicroelectronics.pdf, 2024.
- [38] VL53L8CX - Low-power high-performance 8x8 multizone Time-of-Flight sensor (ToF). <https://www.st.com/en/imaging-and-photonics-solutions/vl53l8cx.html>.
- [39] P-NUCLEO-53L8A1 - STM32 Nucleo pack with X-NUCLEO-53L8A1 expansion board and NUCLEO-F401RE development board. <https://www.st.com/en/ecosystems/p-nucleo-53l8a1.html>.
- [40] X-CUBE-TOP1 - Time-of-Flight sensors software expansion for STM32Cube. <https://www.st.com/en/ecosystems/x-cube-tof1.html>.
- [41] Datasheet - VL53L8CX - Low-power high-performance 8x8 multizone Time-of-Flight (ToF) sensor. <https://www.st.com/resource/en/datasheet/vl53l8cx.pdf>.
- [42] High Voltage Power Monitor | Monsoon Solutions. <https://www.monsoon.com/high-voltage-power-monitor>.
- [43] Photodiode, LSSPD-6-2P-0. <http://www.lightsensing.com/product375.html>.
- [44] Operational Amplifier, OPA320AIDBVR. <https://www.ti.com/product/OPA320/part-details/OPA320AIDBVR>.
- [45] Diode, SMS7630-079LF. <https://www.skyworksinc.com/Products/Diodes/SMS7630-Series>.
- [46] Comparator, TLV7031DBVR. <https://www.ti.com/product/TLV7031/part-details/TLV7031DBVR>.
- [47] Advanced Design System (ADS)s. <https://www.keysight.com/us/en/products/software/pathwave-design-software/pathwave-advanced-design-system.html>.
- [48] Coin Battery CR2016. <https://www.digikey.com/en/products/detail/panasonic-bsg/CR2016/6681>.
- [49] Ultra Low Power Operational Amplifier, MAX9917. <https://www.analog.com/media/en/technical-documentation/data-sheets/max9914-max9917.pdf>.
- [50] Camera Calibration. https://docs.opencv.org/4.x/dc/dbb/tutorial_py_calibration.html.
- [51] pyapriltags: Python bindings for the Apriltags library. <https://pypi.org/project/pyapriltags/>.
- [52] pyzbar - PyPI. <https://pypi.org/project/pyzbar/>.
- [53] VL53L4CX - Time-of-Flight (ToF) sensor with extended range measurement. <https://www.st.com/en/imaging-and-photonics-solutions/vl53l4cx.html>.
- [54] myCobot - Elephant Robotics. <https://www.elephantrobotics.com/en/mycobot-en/>.
- [55] Joshua Springer, Gylfi Þór Guðmundsson, and Marcel Kyas. A precision drone landing system using visual and ir fiducial markers and a multi-payload camera. *arXiv preprint:2403.03806*, 2024.
- [56] Kang Min Bae, Hankyeol Moon, and Song Min Kim. Supersight: Sub-cm nlos localization for mmwave backscatter. In *ACM MobiSys*, 2024.
- [57] Kang Min Bae, Hankyeol Moon, Sung-Min Sohn, and Song Min Kim. Hawkeye: Hectometer-range subcentimeter localization for large-scale mmwave backscatter. In *ACM MobiSys*, 2023.
- [58] Mohamed I Hany, Hamada Rizk, and Moustafa Youssef. Airtags for human localization, not just objects. In *ACM SIGSPATIAL International Workshop on Geo-Privacy and Data Utility for Smart Societies*, 2024.
- [59] Minghui Zhao, Tyler Chang, Aditya Arun, Roshan Ayyalasomayajula, Chi Zhang, and Dinesh Bharadia. Uloc: Low-power, scalable and cm-accurate ubw-tag localization and tracking for indoor applications. *ACM IMWUT*, 2021.
- [60] Roshan Ayyalasomayajula, Deepak Vasisht, and Dinesh Bharadia. Bloc: Csi-based accurate localization for ble tags. In *ACM CoNEXT*, 2018.
- [61] Bo Liang, Purui Wang, Renjie Zhao, Heyu Guo, Pengyu Zhang, Junchen Guo, Shunmin Zhu, Hongqiang Harry Liu, Xinyu Zhang, and Chenren Xu. RF-Chord: Towards deployable RFID localization system for logistic networks. In *USENIX NSDI*, 2023.
- [62] Dianhan Xie, Xudong Wang, and Aimin Tang. Metasight: localizing blocked rfid objects by modulating nlos signals via metasurfaces. In *ACM MobiSys*, 2022.

- [63] Joseph DeGol, Timothy Bretl, and Derek Hoiem. Chromatag: A colored marker and fast detection algorithm. In *IEEE ICCV*, 2017.
- [64] Burak Benligiray, Cihan Topal, and Cuneyt Akinlar. Stag: A stable fiducial marker system. *Image and Vision Computing*, 2019.
- [65] Vincent Lepetit et al. Monocular model-based 3d tracking of rigid objects: A survey. *Foundations and Trends in Computer Graphics and Vision*, 2005.
- [66] BRK Mantha and B Garcia de Soto. Designing a reliable fiducial marker network for autonomous indoor robot navigation. In *International Symposium on Automation and Robotics in Construction*, 2019.
- [67] Yibo Liu, Hunter Schofield, and Jinjun Shan. Navigation of a self-driving vehicle using one fiducial marker. In *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, 2021.
- [68] Phong Ha Nguyen, Muhammad Arsalan, Ja Hyung Koo, et al. Lightdenseyolo: A fast and accurate marker tracker for autonomous uav landing by visible light camera sensor on drone. *Sensors*, 2018.
- [69] Mahathi Bhargavapuri, Animesh Kumar Shastry, Harsh Sinha, Soumya Ranjan Sahoo, and Mangal Kothari. Vision-based autonomous tracking and landing of a fully-actuated rotorcraft. *Control Engineering Practice*, 2019.
- [70] Xu Zhang, Mingyang Li, Jian Hui Lim, Yiwei Weng, Yi Wei Daniel Tay, Hung Pham, and Quang-Cuong Pham. Large-scale 3d printing by a team of mobile robots. *Automation in Construction*, 2018.
- [71] Masanao Koeda, Daiki Yano, Naoki Shintaku, Katsuhiko Onishi, and Hiroshi Noborio. Development of wireless surgical knife attachment with proximity indicators using aruco marker. In *Human-Computer Interaction. Interaction in Context*, 2018.
- [72] Gustavo Alarcón-Nieto, Jacob M Graving, James A Klarevas-Irby, et al. An automated barcode tracking system for behavioural studies in birds. *Methods in Ecology and Evolution*, 2018.
- [73] Tzu-Hsu Yu and Hsin-Mu Tsai. Remark: Privacy-preserving fiducial marker system via single-pixel imaging. In *ACM MobiCom*, 2023.
- [74] Hao Pan, Feitong Tan, Yi-Chao Chen, Gaoang Huang, Qingyang Li, Wenhao Li, Guangtao Xue, Lili Qiu, and Xiaoyu Ji. Docam: depth sensing with an optical image stabilization supported rgb camera. In *ACM MobiCom*, 2022.
- [75] Jinrui Zhang, Huan Yang, Ju Ren, Deyu Zhang, Bangwen He, Ting Cao, Yuanchun Li, Yaoxue Zhang, and Yunxin Liu. Mobidepth: Real-time depth estimation using on-device dual cameras. In *ACM MobiCom*, 2022.
- [76] Simeng Qiu, Hadi Amata, and Wolfgang Heidrich. Moirétag: Angular measurement and tracking with a passive marker. In *ACM SIGGRAPH*, 2023.
- [77] Jingyi Ning, Lei Xie, Zhihao Yan, Yanling Bu, and Jun Luo. Moirévision: A generalized moiré-based mechanism for 6-dof motion sensing. In *ACM MobiCom*, 2024.
- [78] Pengjin Xie, Lingkun Li, Jiliang Wang, and Yunhao Liu. Litag: localization and posture estimation with passive visible light tags. In *ACM SenSys*, 2020.
- [79] Pengjin Xie, Lingkun Li, Jiliang Wang, and Yunhao Liu. Passive visible light tag system for localization and posture estimation. *IEEE TMC*, 2024.
- [80] Xiao Zhang, Hanqing Guo, James Mariani, and Li Xiao. U-star: An underwater navigation system based on passive 3d optical identification tags. In *ACM MoobiCom*, 2022.
- [81] Chi Lin, Jianan Lin, Jie Xiong, Qiwei Wang, Lei Wang, Guowei Wu, Xin Fan, and Zhongxuan Luo. Uwbeacon: Lighting up centimeter-level underwater positioning. In *ACM MobiCom*, 2024.
- [82] Lei Yang, Yekui Chen, Xiang-Yang Li, Chaowei Xiao, Mo Li, and Yunhao Liu. Tagoram: Real-time tracking of mobile rfid tags to high precision using cots devices. In *ACM MobiCom*, 2014.
- [83] Longfei Shangguan and Kyle Jamieson. The design and implementation of a mobile rfid tag sorting robot. In *Proceedings of the 14th annual international conference on mobile systems, applications, and services*, pages 31–42, 2016.
- [84] Laura Dodds, Isaac Perper, Aline Eid, and Fadel Adib. A handheld fine-grained rfid localization system with complex-controlled polarization. In *ACM MobiCom*, 2023.
- [85] Impinj xSpan and xArray Gateways. https://support.impinj.com/hc/article_attachments/31811120112147f, 2024.
- [86] AirTag. <https://www.apple.com/airtag/>, 2024.
- [87] Pengyu Zhang, Dinesh Bharadia, Kiran Joshi, and Sachin Katti. Hitchhike: Practical backscatter using commodity wifi. In *ACM SenSys*, 2016.
- [88] Zicheng Chi, Xin Liu, Wei Wang, Yao Yao, and Ting Zhu. Leveraging ambient lte traffic for ubiquitous passive communication. In *ACM SIGCOMM*, 2020.
- [89] Vamsi Talla, Mehrdad Hesar, Bryce Kellogg, Ali Najafi, Joshua R Smith, and Shyamnath Gollakota. Lora backscatter: Enabling the vision of ubiquitous connectivity. *ACM IMWUT*, 2017.
- [90] Kang Min Bae, Namjo Ahn, Yoon Chae, Parth Pathak, Sung-Min Sohn, and Song Min Kim. Omnisscatter: extreme sensitivity mmwave backscattering using commodity fmcw radar. In *ACM MobiSys*, 2022.
- [91] Yoon Chae, Zhenzhe Lin, Kang Min Bae, Song Min Kim, and Parth Pathak. mmComb: High-speed mmWave commodity WiFi backscatter. In *USENIX NSDI*, 2024.
- [92] Farzan Dehbashi, Ali Abedi, Tim Brecht, and Omid Abari. Verification: can wifi backscatter replace rfid? In *ACM MobiCom*, 2021.
- [93] Jiangtao Li, Angli Liu, Guobin Shen, Liquan Li, Chao Sun, and Feng Zhao. Retro-vlc: Enabling battery-free duplex visible light communication for mobile and iot applications. In *ACM HotMobile*, 2015.
- [94] Sihua Shao, Abdallah Khreishah, and Hany Elgala. Pixelated vlc-backscattering for self-charging indoor iot devices. *IEEE Photonics Technology Letters*, 2017.
- [95] Yue Wu, Purui Wang, Kenuo Xu, Lilei Feng, and Chenren Xu. Turboboosting visible light backscatter communication. In *ACM SIGCOMM*, 2020.
- [96] Kenuo Xu, Chen Gong, Bo Liang, Yue Wu, Boya Di, Lingyang Song, and Chenren Xu. Low-latency visible light backscatter networking with RetroMUMIMO. In *ACM SenSys*, 2022.
- [97] Rens Bloom, Marco Zúñiga Zamalloa, and Chaitra Pai. Luxlink: creating a wireless link from ambient light. In *ACM SenSys*, 2019.
- [98] Seyed Keyarash Ghiasi, Marco A. Zúñiga Zamalloa, and Koen Langendoen. A principled design for passive light communication. In *ACM MobiCom*, 2021.
- [99] Seyed Keyarash Ghiasi, Vivian Dsouza, Koen Langendoen, and Marco Zúñiga. Spectralux: Towards exploiting the full spectrum with passive vlc. In *ACM/IEEE IPSN*, 2023.
- [100] Yu-Lin Wei, Chang-Jung Huang, Hsin-Mu Tsai, and Kate Ching-Ju Lin. Celli: Indoor positioning using polarized sweeping light beams. In *ACM MobiSys*, 2017.
- [101] Sihua Shao, Abdallah Khreishah, and Issa Khalil. Retro: Retroreflector based visible light indoor localization for real-time tracking of iot devices. In *IEEE INFOCOM*, 2018.
- [102] Sihua Shao, Abdallah Khreishah, and Juan Paez. Passiveretro: Enabling completely passive visible light localization for iot applications. In *IEEE INFOCOM*, 2019.
- [103] Wei Li, Tuochao Chen, Zhe Ou, Xin Wen, Zichen Xu, and Chenren Xu. RetroFlex: enabling intuitive human-robot collaboration with flexible retroreflective tags. *CCF Transactions on Pervasive Computing and Interaction*, 2022.
- [104] Miguel Chávez Tapia, Talia Xu, Zehang Wu, and Marco Zúñiga Zamalloa. Sun-box: Screen-to-camera communication with ambient light. *ACM IMWUT*, 2022.
- [105] Miguel A Chávez Tapia, Talia Xu, and Marco Zúñiga Zamalloa. Sol-fi: Enabling joint illumination and communication in enclosed areas with sunlight. In *ACM/IEEE IPSN*, 2024.
- [106] Purui Wang, Lilei Feng, Guojun Chen, et al. Renovating road signs for infrastructure-to-vehicle networking: A visible light backscatter communication and networking approach. In *ACM MobiCom*, 2020.
- [107] Chenren Xu, Kenuo Xu, Lilei Feng, and Bo Liang. Retrov2x: A new vehicle-to-everything (v2x) paradigm with visible light backscatter networking. *Fundamental Research*, 2023.
- [108] Kenuo Xu, Kexing Zhou, Chengxuan Zhu, Shanghang Zhang, Boxin Shi, Xiaoqiang Li, Tiejun Huang, and Chenren Xu. When visible light (backscatter) communication meets neuromorphic cameras in v2x. In *ACM HotMobile*, 2023.
- [109] Fangzhou Mu, Carter Sifferman, Sacha Jungerman, Yiquan Li, Mark Han, Michael Gleicher, Mohit Gupta, and Yin Li. Towards 3d vision with low-cost single-photon cameras. In *IEEE/CVF CVPR*, 2024.
- [110] Gregor Luetzenburg, Aart Kroon, and Anders A Bjørk. Evaluation of the apple iphone 12 pro lidar for an application in geosciences. *Scientific reports*, 2021.
- [111] Stefano Tavani, Andrea Billi, Amerigo Corradetti, Marco Mercuri, Alessandro Bosman, Marco Cuffaro, Thomas Seers, and Eugenio Carminati. Smartphone assisted fieldwork: Towards the digital transition of geoscience fieldwork using lidar-equipped iphones. *Earth-Science Reviews*, 2022.
- [112] Jiawei Hu, Yanxiang Wang, Hong Jia, Cheng Jiang, Mahbub Hassan, Brano Kusy, and Wen Hu. Lidarspectra: Synthetic indoor spectral mapping with low-cost lidars. In *ACM/IEEE IPSN*, 2024.
- [113] Ruichun Ma, Shicheng Zheng, Hao Pan, Lili Qiu, Xingyu Chen, Liangyu Liu, Yihong Liu, Wenjun Hu, and Ju Ren. Automs: Automated service for mmwave coverage optimization using low-cost metasurfaces. In *ACM MobiCom*, 2024.
- [114] Alice Ruget, Max Tyler, Germán Mora Martín, Stirling Scholes, Feng Zhu, Istvan Gyongy, Brent Hearn, Steve McLaughlin, Abderrahim Halimi, and Jonathan Leach. Pixels2pose: Super-resolution time-of-flight imaging for 3d pose estimation. *Science Advances*, 2022.
- [115] Jan Pleterski, Gašper Škulj, Corentin Esnault, Jernej Puc, Primož Podržaj, et al. Miniature mobile robot detection using an ultra-low resolution time-of-flight sensor. *IEEE Transactions on Instrumentation and Measurement*, 2023.
- [116] Clara Callenberg, Zheng Shi, Felix Heide, and Matthias B Hullin. Low-cost spad sensing for non-line-of-sight tracking, material classification and depth imaging. *ACM TOG*, 2021.
- [117] Cienna N Becker and Lucas J Koerner. Plastic classification using optical parameter features measured with the tmf8801 direct time-of-flight depth sensor. *Sensors*, 2023.
- [118] Carter Sifferman, Yeping Wang, Mohit Gupta, and Michael Gleicher. Unlocking the performance of proximity sensors by utilizing transient histograms. *IEEE RA-L*, 2023.
- [119] Sacha Jungerman, Atul Ingle, Yin Li, and Mohit Gupta. 3d scene inference from transient histograms. In *ECCV*, 2022.
- [120] Justin Chan, Ananditha Raghunath, Kelly E Michaelsen, and Shyamnath Gollakota. Testing a drop of liquid using smartphone lidar. *ACM IMWUT*, 2022.