



Catch Me If You Can: Laser Tethering with Highly Mobile Targets

Charles J. Carver, Hadleigh Schwartz, and Qijia Shao, *Columbia University*;
Nicholas Shade, Joseph Lazzaro, Xiaoxin Wang, Jifeng Liu, and Eric Fossum,
Dartmouth College; Xia Zhou, *Columbia University*

<https://www.usenix.org/conference/nsdi24/presentation/carver>

This paper is included in the
Proceedings of the 21st USENIX Symposium on
Networked Systems Design and Implementation.

April 16–18, 2024 • Santa Clara, CA, USA

978-1-939133-39-7

Open access to the Proceedings of the
21st USENIX Symposium on Networked
Systems Design and Implementation
is sponsored by



Catch Me If You Can: Laser Tethering with Highly Mobile Targets

Charles J. Carver^{†1}, Hadleigh Schwartz^{†1}, Qijia Shao¹, Nicholas Shade², Joseph Lazzaro², Xiaoxin Wang², Jifeng Liu², Eric Fossum², and Xia Zhou¹

[†]Co-primary authors

¹Department of Computer Science, Columbia University, ²Thayer School of Engineering, Dartmouth College

{cjc, hadleigh, qijia, xia}@cs.columbia.edu,

{nicholas.r.shade.th, joseph.p.lazzaro.th, xiaoxin.wang, jifeng.liu, eric.r.fossum}@dartmouth.edu

Abstract

Conventional wisdom holds that laser-based systems cannot handle high mobility due to the strong directionality of laser light. We challenge this belief by presenting Lasertag, a *generic framework* that tightly integrates laser steering with optical tracking to maintain laser connectivity with high-velocity targets. Lasertag creates a constantly connected, laser-based tether between the Lasertag core unit and a remote target, irrespective of the target's movement. Key elements of Lasertag include (1) a novel optical design that superimposes the optical paths of a steerable laser beam and image sensor, (2) a lightweight optical tracking mechanism for passive retroreflective markers, (3) an automated mapping method to translate scene points to laser steering commands, and (4) a predictive steering algorithm that overcomes limited image sensor frame rates and laser steering delays to quadruple the steering rate up to 151 Hz. Experiments with the Lasertag prototype demonstrate that, on average, Lasertag transmits a median 0.97 of laser energy with a median alignment offset of only 1.03 cm for mobile targets accelerating up to 49 m/s², with speeds up to 6.5 m/s, and distances up to 6 m ($\approx 47^\circ$ /s). Additional experiments translate the above performance to a 10^{-8} median bit error rate across trials when transmitting a 1 Gbps, on-off keying signal. Lasertag paves the way for various laser applications (e.g., communication, sensing, power delivery) in mobile settings. A demonstration video of Lasertag is available at:

mobilex.cs.columbia.edu/lasertag

1 Introduction

The physical properties of laser light make it an excellent medium for numerous applications. Examples include high-bandwidth communication due to its fast modulation speeds [25, 39, 53], efficient power delivery thanks to its high energy density [27, 54, 55, 58], and fine-grained sensing given its nanometer-level spectral widths [19, 43, 112, 149].

Despite its potential, the inherent directionality of laser light has precluded its use in highly-mobile settings. Although existing free-space optics (FSO) systems can track and steer to fast-moving objects with *predetermined* trajectories (e.g., satellites) at *kilometer-level* distances, these systems are unsuitable for emerging mobile applications – such as virtual reality (VR) streaming and mobile power delivery – that exhibit *arbitrary* trajectories at *meter-level* distances. Fundamentally, this discrepancy is due to the higher angular tracking rate required for near-range targets. For example, a target 1m away traveling at 5m/s requires an angular tracking rate of




Figure 1: Lasertag maintains constant laser alignment with a flying drone equipped with a passive retroreflective marker.

approximately 79° /s, in contrast with the 1.5° /s [141] rate required to maintain laser connectivity with a low earth orbit satellite moving at 7.8 km/s [96]. Compounding FSO's insufficient tracking rate for emerging applications, these systems are often bulky and unsuitable for mobile environments.

In this work, we bridge the gap between traditional FSO systems and emerging mobile applications by enabling fast tracking and laser steering using portable, off-the-shelf hardware. To accomplish this task, we first consider various optical designs and their corresponding characteristics. While diffusing the beam can mitigate the alignment issue, it inevitably sacrifices supporting range and requires high-power laser diodes, resulting in low energy efficiency. An alternative is to scan the narrow laser beam in search of the target [16, 69, 114], which entails delays of hundreds of milliseconds [18]. Such delays are problematic for tracking fast motions such as consumer drone and human head movements (both multiple m/s [28, 118]). Additionally, once the target is acquired, additional scans may be required to maintain tracking [16, 59, 83, 86, 116] resulting in breaks in the laser link, which is unsuitable for applications requiring constant connections (e.g., data transmission).

A more efficient approach is to *separately* track the target's movement and then steer the laser beam directly to the target's location. While object tracking and laser steering are well-explored endeavours on their own, their integration is nontrivial on multiple fronts. First, the narrow-beam nature of laser light leads to a low tolerance for positioning errors, rendering existing localization technologies insufficient for accurate laser steering. Second, even if the target is perfectly positioned, it is challenging to translate the object's 3D location to the 2D reference frame used by the laser-steering device: any offset between the positioning and laser-steering device causes depth

ambiguity (i.e., different object depths necessitate different steering angles). This offset often cannot be measured due to the unknown origin points within the positioning and steering devices. Finally, mapping the steering device's input to an outgoing beam angle requires exhaustive measurements which are susceptible to human errors. This is further complicated by additional optics (e.g., wide-angle lenses) used in portable laser-steering systems [16–18, 75, 84], which affects both the tracking unit's perception of the scene and outgoing laser angle.

We present Lasertag to address the above practical issues. Lasertag is a *reconfigurable, application-agnostic platform* that tightly integrates laser steering and optical tracking in support of numerous high mobility applications. As shown in Figure 1, Lasertag provides a laser-based tether between a laser diode and an arbitrary object, irrespective of the object's movement. Leveraging a novel optical design that superimposes the optical paths of a steerable laser beam and single image sensor, we efficiently fuse the power of computer vision with the communication and sensing benefits of laser light. The optical path sharing addresses the issue of depth ambiguity and optical-element complications. Furthermore, we exploit the physics of fluorescence to enable an automated, precise mapping between any point in the scene and the laser steering's input drive signal. Finally, to overcome the limited frame rate of low-cost image sensors and non-negligible delays of laser steering devices, we propose a predictive steering algorithm that forecasts the target's future location, interpolates its intermediate locations, and proactively steers to the interpolated points until a new sensor reading is ready.

We fabricate a Lasertag prototype using off-the-shelf hardware, and evaluate its performance in the context of three applications where unidirectional,¹ high-mobility laser tethering is a necessity: ground-to-drone communication, laser-based streaming of VR content, and mobile power delivery. Since application-specific performance fundamentally relies on the ability to maintain a laser connection with a moving target, we focus on evaluating Lasertag's *tethering efficacy* by measuring the offset between the laser beam's center and the marker, as well as the received laser power at the target. We then provide additional experiments and analyses to translate the measured tethering performance to application-specific performance. Finally, we examine Lasertag's individual components, its robustness, and power consumption. In summary:

- Lasertag, on average, tethers a laser beam to a mobile target's light sensor with a median offset of only 1.03 cm, delivering a median normalized laser power of 0.97 for targets accelerating up to 49 m/s^2 , traveling with speeds up to 6.5 m/s , and at distances up to 6 m ($\approx 47^\circ/\text{s}$).
- The predictive steering algorithm boosts the tracking and steering rate from 36 Hz to 151 Hz, delivering 15% higher laser power and reducing alignment offset by 12% on

¹ Bidirectional applications can be supported with either a second Lasertag unit on the mobile node or alternative communication techniques.




Figure 2: Challenge of depth ambiguity, where changes in an object's depth (z_1, z_2, z_3) require different steering angles ($\beta_1, \beta_2, \beta_3$).

average across all velocity and trajectory patterns.

- Translating tethering efficacy to communication performance, Lasertag maintains a median 10^{-8} bit error rate (BER) for a 1 Gbps on-off keying (OOK) signal using predictive steering, a 54 times smaller BER than baseline steering.
- Lasertag's optical tracking reliably detects a passive retroreflective marker up to 6 m in the presence of confounding background objects (e.g., LEDs, reflective objects) and strong ambient light (up to 200 kLux).
- Lasertag supports a $\pm 95^\circ$ steering range with ≤ 3 dB loss.

2 System Challenges

The integration of laser steering and object tracking faces the following challenges.

Small Beam Divergence. Laser-based systems typically utilize narrowly collimated beams with small divergences ($0.0005^\circ - 0.005^\circ$) [66] for enhanced range and energy efficiency. This results in low tolerance of localization/tracking errors. Existing localization methodologies using radio frequencies (e.g., GPS [8], Wi-Fi [68, 120], and Bluetooth [3, 148]) have localization errors ranging from tens of centimeters [3, 68, 120, 148] to multiple meters (outdoor GPS [135]). Similarly, for systems relying on inertial measurement units, positioning errors can accumulate to meters in only a short duration [36, 72]. These localization errors can easily cause a laser beam to miss a target at meter-level distances, even with perfect steering control. Diffused laser beams may mitigate this issue, but at the cost of significantly reduced energy efficiency.

Depth Ambiguity. Even with perfect localization, it is challenging to translate the localization's 3D reference frame to the 2D reference frame of the outgoing laser beam. All common laser-steering devices, from bulky mechanical gimbals to micro-electromechanical (MEMS) mirrors, convert an input signal to a mechanical deflection around a mirror's center point. This 2D mechanical deflection steers the laser beam in 3D space so that it reaches the first object lying in its path.

As shown in Figure 2, the offset d between the positioning mechanism and laser-steering device can cause depth ambiguity when determining the 2D angle for steering the laser. This offset leads to different views for the two units. Hence, even locations that differ only in their depths (i.e., z_0, z_1, z_2) require different steering angles ($\beta_0, \beta_1, \beta_2$) for the laser beam to reach them. Therefore, to calculate the correct outgoing 2D angle that will reach a target object in 3D space, an accurate measurement of the offset d is crucial. Obtaining this measurement

is often impractical, as the positioning and steering devices typically have unknown origin points somewhere within their packaging, and even mm-level measurement inaccuracies would cause the steered laser beam to miss its target.

Mapping Overhead. Compounding the above challenges, the laser steering mechanism must be precisely characterized to map input drive signals to outgoing steering angles, which becomes even more difficult when using additional lenses.

Traditional free-space optics systems utilize bulky, slow, and expensive gimbals to steer laser light to the target [115]. A major benefit of these systems is their simple mapping between input signal and output angle [132], often characterized by the manufacturer during fabrication. High mobility applications require smaller and faster steering devices, typically in the form of MEMS mirrors. Unfortunately, MEMS mirrors exhibit model-dependent nonlinear relationships between input signal and output signal [89]. Mapping MEMS mirror drive signals to deflection angles thus requires an independent characterization, which is prone to human error, time-consuming, and reduces the accuracy of the system.

Worsening the above issue, fast-steering MEMS mirrors have an extremely limited steering range (e.g., $\pm 7^\circ$ mechanical tilt) and therefore require additional optics to expand the steering range to a desirable field-of-view (FOV) [16, 18]. These secondary optics not only increase the mapping complexity and introduce more opportunities for error, but also change the outgoing ray geometry so they no longer exit the optical system from a single point. Instead, these rays exit the optical system at different spatial positions along the exit lens depending on steering angle, such that the offset d (Figure 2) between the laser steering system and object positioning system becomes angle-dependent. This significantly increases the colocation complexity and requires an additional mapping.

3 Lasertag Design

To address the above challenges, we present Lasertag, a generic framework that fuses laser steering with optical tracking. We choose optical tracking because it provides the most flexibility and ease of colocation compared to alternatives. We now elaborate on the four key design elements of Lasertag.

3.1 Efficient Optical Path Sharing

The first design element tackles the issue of translating the localization reference frame to the laser steering reference frame. We propose a novel optical design that intrinsically fuses the two together, efficiently sharing the optical path between the outgoing laser light and a single image sensor. This optical arrangement eliminates manual and time-consuming measurements to enable 2D laser steering in 3D space.

The proposed optical design requires only a single image sensor for tracking. To understand the rationale, we envision a hypothetical optical circuit consisting of an image sensor and single focusing lens. The lens focuses light from the scene onto the image sensor, creating a 2D projection of the 3D world. For each patch of focused light, a chief ray (CR) exists, originating




Figure 3: The optical design of Lasertag, enabling outgoing laser light (red) to be mapped 1:1 with incoming scene light (purple).

from a 3D scene point and passing through the lens to project onto a 2D pixel patch. With this optical circuit in mind, we can then consider adding an imaginary laser steering mechanism that is perfectly collocated with the image sensor. This laser steering device is capable of steering a narrow laser beam (LR) out through the same focusing lens, retracing the steps of the incoming CR to reach the same object as imaged by the sensor. In this way, a 2D pixel is mapped to a 2D angle relative to the image sensor, thereby eliminating depth ambiguity.

To realize this optical design, we exploit two intrinsic properties of laser light: strong linear polarization and narrow emission wavelength. Specifically, we overlay the laser’s outgoing optical path with an image sensor’s incoming optical path (CR + LR) using a beamsplitter, which reflects laser light out of the system and transmits incoming scene light onto the image sensor. Our methodology – inspired by the field of microscopy [97, 109] – maximizes the steered laser energy without sacrificing imaging quality, thus enabling efficient laser tethering. Figure 3 illustrates Lasertag’s optical design, where additional lenses (e.g., triplet lens) create the proper ray orientation for the optical circuit. The final fisheye lens enables full-hemisphere steering and imaging of the scene.

To support our goal, we consider beamsplitters that redirect light using either fixed ratios, wavelength (i.e., dichroic beamsplitters), or polarization. Fixed-ratio beamsplitters allow only a fixed ratio of transmission to reflection (e.g., 10:90, 50:50), sacrificing transmitted laser power for imaging brightness (or vice versa). Given the narrow spectral width of laser light, dichroic beamsplitters, which reflect light above/below a certain wavelength and transmit the remaining, seem an obvious choice. Unfortunately, dichroic beamsplitters cause the imaged scene to be extremely skewed in color (since only light above or below the beamsplitter threshold would reach the image sensor), and are often manufactured on thin plates resulting in poor image quality due to ghosting [127].

To solve these challenges, we design our system around a laser-line polarizing beamsplitter. These beamsplitters reflect linearly polarized light lying within a narrow wavelength range, and transmit light that is either (a) polarized in the other direction, or (b) at a wavelength above the polarizing range. This design leverages the fact that laser light is intrinsically highly polarized and narrow in wavelength [108]. As a result, our design can operate in two modes: one in which true




Figure 4: Workflow for isolating the passive marker from cluttered backgrounds and steering the laser to the marker's center.

scene color is desired, and one in which maximum scene brightness is required. In the first case, since ambient light is typically unpolarized, the image sensor sees all wavelengths with approximately 50% attenuation in overall brightness. In the second case, artificial light is injected into the scene at a wavelength above the beamsplitter's polarizing wavelength range, at which it transmits all light, regardless of polarization orientation. This enables the bulk of sensing light to be transmitted back to the image sensor. Our prototype (§4) implements the second mode. Notably, this requires the use of *separate laser and sensing wavelengths*: a laser wavelength *inside* the beamsplitter's specified polarizing wavelength range, and a sensing wavelength *outside* it. Lasertag can be deployed with any pair of wavelengths that satisfy this basic constraint.² Since commercial beamsplitters are available in a variety of polarizing ranges, the appropriate choice of laser and sensing wavelengths primarily depends on application requirements and channel characteristics. For instance, Lasertag should be equipped with blue-green light sources for underwater scenarios, given the lower attenuation of these wavelengths in water. In contrast, for aerial power delivery applications, it would be favorable to use an IR laser to match the responsivity of typical solar cells. Additionally, our design imposes no restrictions on the light sources' optical powers, so that safety, power consumption, and signal-to-noise ratio (SNR) requirements can also be considered. Finally, since the outgoing light is completely separable from the incoming light, the input laser parameters have no impact on Lasertag's sensing performance.

While several elements of the design are explored in existing works (e.g., MEMS mirrors for laser beam steering [16, 75], beamsplitting optics for colocation of laser steering and detection components [6, 23, 32], fisheye lenses for expanding laser steering range [16, 18, 22, 151]), the integration of these elements for tracking and steering is nontrivial. The optical characteristics must be carefully designed to create the appropriate ray geometry for imaging and steering.

3.2 Fast Tracking with Retroreflective Imaging

The second design element identifies the target object from each image frame. Object detection is well studied in computer vision (CV). The proposed marker design and algorithm are built upon prior CV techniques and inspired by existing works leveraging retroreflectors for tracking [83, 137, 147]. The algorithm incorporates simple heuristics that differentiate

the passive marker from interfering scene objects and run real-time tracking on resource-constrained hardware.

Retroreflective Marker. The target object is equipped with a passive marker made of retroreflective tape, cut into a circle, with a small hole in the center to host a light sensor. After adding a wide-angle LED to Lasertag, the scene is illuminated with a specific wavelength that retroreflects off the marker and back to the image sensor. In general, these retroreflections will be significantly brighter than other objects in the scene, so the target will appear as a blob that can be segmented from background objects. The laser light is then steered to the center of the blob, so that it can pass through the hole to the sensor (e.g., photodiode for receiving data, solar-cell for harvesting energy). Lasertag's optical design ensures that the laser beam is sufficiently narrow so the majority of light passes through.

Although retroreflections are generally brighter than arbitrary scene objects, active luminaries emitting at the same wavelength have the potential to be brighter. To overcome this challenge, we leverage the fact that since light is not retroreflected by the center of the marker (due to the cutout), the marker's center pixels will be dimmer than the remainder of the marker. Confounding luminaries in the scene, however, will have a roughly uniform illumination within the blob (e.g., LEDs will emit light from the center of their package). Consequently, we identify the marker by taking the ratio of the blob's overall brightness to the blob's center brightness, favoring blobs with higher ratios. We add an additional weight for the roundness of the blob, since the circular marker will remain elliptical even under rotation. The entire image processing algorithm is described in the following section.

Marker Detection and Tracking. To achieve high marker tracking rates, we develop a lightweight marker detection algorithm based on traditional CV blob detection [26, 45] and implement *trajectory-based region of interest (ROI) windowing* to increase both frame capture and detection speeds. The key insight of our detection algorithm is that the passive marker registers in the image as a distinctly bright region with a known shape. Based on these characteristics, we isolate the marker in each frame and determine the pixel corresponding to its center. This process is visually described in Figure 4. We note that regardless of the marker's angular orientation, its center remains constant in the 2D frames. In addition to localizing the marker, the output of our detection algorithm informs our trajectory-based ROI windowing algorithm (Algorithm 1) to increase Lasertag's tracking rate. ROI windowing is a camera

²Advanced Lasertag designs can consider multispectral beam splitters to support multiple input laser wavelengths.




Figure 5: (Left) Fluorescent cap design enabling automatic pixel-to-voltage mapping. (Right) After steering the laser to each white dot, we fit a polynomial surface to obtain our mapping functions.

feature that enables reading a smaller window of pixel values, effectively zooming in to a region of interest. This reduces the number of pixel values the camera must process and thus decreases image capture time [92], which can facilitate real-time tracking of objects moving at high speeds [111]. We set an ROI window that hones in on the marker and follows its trajectory.

As shown in Figure 4, after capturing each frame, we binarize the image using a threshold automatically determined with Otsu’s method [9]. The retroreflective marker and any confounding luminaries then appear in the image as clusters of connected white pixels (i.e., blobs). We run blob detection on the image to determine the center pixel, bounding box, and roundness for each region. To isolate our passive marker, we assign each blob a score based on our marker heuristics (favoring round regions with high ratio of overall brightness to center brightness). We then save the blob with the highest score and return its center pixel coordinate. Finally, we update the image sensor registers to move the ROI window based on the updated marker location. If no marker is detected, the ROI is expanded to encompass the whole image sensor so that the target can be re-acquired. Our algorithm’s simplicity, afforded by its tight coupling with the physical design of our passive marker, enables localization and tracking within milliseconds.

Notably, Lasertag’s framework is inherently flexible and works with any conventional CV detection technique. While we demonstrate our system with a passive circular retroreflective marker, Lasertag is compatible with markers of any shape and size (e.g., a 3D sphere). Furthermore, Lasertag can support markerless tracking and active feedback as detailed in §6.

3.3 Fluorescence-based Optical Mapping

The third design element maps the marker’s position on the image sensor to the steering command required to steer the laser to the object. Lasertag performs beam steering with a MEMS mirror – a compact, electromagnetically driven device in which deflection angles are determined by input voltages. We enable automatic pixel-to-voltage mapping via a short, one-time calibration process that is independent of the environment. The calibration process leverages fluorescence – a physical phenomenon in which a substance is stimulated by the light it absorbs, causing it to emit light of a different color – to shift the outgoing laser’s wavelength to one always visible to the image sensor. With our optical design, this wavelength shift allows us to leverage the embedded image sensor to *automatically* map all pixel locations to steering voltages,

Input: μ : Sensor mode, α or τ .

Output: (x,y) : Absolute center pixel of the marker.

```

1 Function Track( $\mu$ ):
2   while True do
3      $I \leftarrow$  Capture new image
4      $(x',y') \leftarrow$  Detect( $I$ )
5     if  $(x',y') = \text{None}$  then
6        $\mu \leftarrow \alpha$ 
7       Set sensor resolution to  $W_\alpha \times H_\alpha$ 
8       Set sensor readout offset to  $(0,0)$ 
9       Set sensor readout area to  $(W,H)$ 
10      Skip to next iteration
11    if  $\mu = \alpha$  then
12       $\mu \leftarrow \tau$ 
13       $t_x \leftarrow (x' - W_\alpha/2) * W / W_\alpha$ 
14       $t_y \leftarrow (y' - H_\alpha/2) * H / H_\alpha$ 
15      Set sensor resolution to  $W_\tau \times H_\tau$ 
16      Set sensor readout offset to  $(t_x, t_y)$ 
17      Set sensor readout area to  $(W/z, H/z)$ 
18    else
19       $o_x \leftarrow (x' - W_\tau/2) * (W/z) / W_\tau$ 
20       $o_y \leftarrow (y' - H_\tau/2) * (H/z) / H_\tau$ 
21       $(t_x, t_y) \leftarrow$  Updated sensor readout offset
22       $(x,y) \leftarrow (o_x + t_x, o_y + t_y)$ 
23      Set sensor readout offset to  $(t_x + o_x, t_y + o_y)$ 
24    return  $(x,y)$ 

```

Algorithm 1: Marker tracking algorithm for an image sensor with a $W \times H$ active area. After marker acquisition ($\mu = \alpha$), the $W_\alpha \times H_\alpha$ resolution is scaled to $W_\tau \times H_\tau$ for tracking ($\mu = \tau$) and the active pixel readout window is reduced to $W/z \times H/z$.

bypassing the requirement of accurate angular measurements.

Specifically, we borrow from the field of fluorescence microscopy [97, 109] and use fluorescent powder to shift the laser wavelength *after* it has been steered to one visible to the image sensor. We design a small plastic cap that is placed over the optical circuit’s exit lens during the calibration process (Figure 5). It is coated with either an up-converting [40] or down-converting phosphor [82]³. The cap is a hollow hemisphere whose interior is a minimum distance of 2 cm from the fisheye lens at all points. This design enforces the minimum required distance between the fluorescent surface and fisheye lens to mitigate comatic aberrations (which result in comet-like tails around the fluorescing spot, leading to mapping errors) [56]. We empirically determine this distance in our system to be 2 cm. To do so, we fix the steering unit’s outgoing laser angle and place a fluorescent surface in the beam’s path in the far field. We then shift the surface towards the unit along the beam until the center coordinate of the fluorescent spot’s blob changes, indicating the minimum distance.

Finally, to generalize the mapping to any pixel on the image, we fit two 3D surfaces with a subset of laser steering voltages and pixel positions. With the fluorescent cap installed, we steer the laser to the center of the fisheye lens, then scan

³For up-converting phosphors, higher wavelength photons are absorbed causing lower wavelength photons to be emitted. Down-converting phosphors absorb lower wavelength photons and emit higher wavelengths.

the beam in an Archimedian spiral [81] to cover N spots on the cap (Figure 5). Steering via the MEMS mirror requires applying two voltages, v_x , and v_y , corresponding to the two orthogonal scanning axes. Given the fisheye lens's large distortion on its edges, we decrease the step size between adjacent steering voltages at larger spiral radii, enabling an accurate mapping despite the spatial compression. Since the sensed pixel position is influenced by both v_x and v_y , we feed (x, y, v_x) and (x, y, v_y) through a multivariate, polynomial regression algorithm to obtain our mapping functions:

$$m_x(x, y) = \sum_{i=0}^3 \sum_{j=0}^3 a_{ij} x^i y^j, \quad m_y(x, y) = \sum_{i=0}^3 \sum_{j=0}^3 b_{ij} x^i y^j,$$

where the coefficients a_{ij} and b_{ij} are determined by the respective regressions. We choose a third-degree polynomial to account for the nonlinear relationship between steering voltage and outgoing beam angle. Overall the calibration takes roughly ten minutes, creating the functions later used by Lasertag to map an identified marker to a steering command in real time. Notably, this streamlined calibration can be automatically run across environments and can be repeated at any point throughout deployment to maintain steering accuracy. This makes it highly flexible compared to alternative systems requiring hard-coded system parameters [4, 126], gimbal positions [20], uplink feedback [1, 64], or complex learning strategies [39].

3.4 Predictive Steering for High Mobility

The fourth design element overcomes the low frame rate of conventional image sensors and delays associated with laser steering devices. Conventional image sensors have a low frame rate (30-60 FPS), which results in a considerable delay between consecutive frames. Additionally, laser steering takes a non-negligible amount of time (on the order of milliseconds). These delays can cause a steering gap between the current and subsequent target locations, particularly when tracking fast-moving objects (e.g., flying drones which accelerate to over 20 m/s² when making sharp turns). As a consequence, laser alignment continuity can easily be disrupted.

To address this problem, we exploit movement continuity on a micro timescale and propose *predictive steering*, wherein we (1) predict the target's expected location in the next frame, (2) interpolate locations from now to the next frame, and (3) steer to interpolated locations until a new frame is ready. This upsampling strategy achieves smoother laser steering and handles fast movement more effectively. While upsampling techniques are common in the domain of audio and image processing [34, 71, 140], they are usually performed offline, requiring the signal to be captured in advance. In the scenario of real-time laser tracking and steering, we must predict the next steering location using only historical data.

The rationale of mobility prediction stems from the continuous nature of motion at micro timescales. Although motion is typically non-deterministic, it is still subject to certain constraints. For example, speed and acceleration are bounded and cannot be changed instantaneously. Therefore, by

Input: M : Prediction buffer size, $\bar{\epsilon}$: Nominal steering delay.

Output: None.

```

1 Function Steer( $M, \bar{\epsilon}$ ):
2   Let  $B, s, b$  be an empty array of length  $M$ 
3   while True do
4      $T \leftarrow$  Time since last iteration
5     if first iteration then
6        $(x_0, y_0) \leftarrow \text{Track}(\alpha)$ 
7     else
8        $(x_0, y_0) \leftarrow \text{Track}(\tau)$ 
9     Append  $(x_0, y_0)$  to  $B$ 
10    Skip to next iteration if  $\text{length}(B) \neq M$ 
11     $(s_x, s_y)^{(0)} \leftarrow B[0]$ 
12     $(b_x, b_y)^{(0)} \leftarrow B[1] - B[0]$ 
13    for  $i \leftarrow 1$  to  $M - 1$  do
14       $(s_x, s_y)^{(i)} \leftarrow \text{Smooth } B[i] \text{ with Eq. (1)}$ 
15       $(b_x, b_y)^{(i)} \leftarrow \text{Estimate trend with Eq. (2)}$ 
16     $(x_T, y_T) \leftarrow (s_x + b_x, s_y + b_y)^{(M-1)}$ 
17     $\Delta \leftarrow$  Time since call to Track(...)
18     $t \leftarrow \Delta + \bar{\epsilon}$ 
19    while  $t \leq T - \bar{\epsilon}$  do
20       $\delta = (y_T - y_0) / (x_T - x_0)$ 
21       $x_t = x_0 + t * (x_T - x_0) / T$ 
22       $y_t = \delta * (x_T - x_0) + y_0$ 
23       $(v_x, v_y) \leftarrow (m_x(x_t, y_t), m_y(x_t, y_t))$ 
24      Set mirror voltage to  $(v_x, v_y)$ 
25       $\epsilon_i \leftarrow$  Time since previous nested iteration
26       $t \leftarrow t + \epsilon_i$ 
27    Delete  $B[0]$ 

```

Algorithm 2: Predictive steering algorithm to continuously steer the laser light to pixels (x_t, y_t) lying between the captured marker position (x_0, y_0) and predicted final location (x_T, y_T) expected at time T . The i^{th} iteration of the double exponential prediction is denoted by a superscript (i) .

utilizing historical data, it is feasible to learn motion dynamics and predict the object's next location. Numerous techniques for motion prediction have been proposed, including double exponential filters [48, 74], Kalman filters [7, 78, 143], and neural networks [44, 113]. In our implementation (§4), we utilize a double exponential predictor due to its low computational overhead. Our framework is generic, however, and can work with any motion prediction model as outlined in §6.

The double exponential predictor is a time-series forecasting method that utilizes a double exponential filter to smooth the input data and estimate future trends [48]. At a high level, the predictor first applies a double exponential filter to the time series data and obtains a smoothed estimate of the underlying trend [98]. Then, it predicts future values by adding the estimated slope of the trend to the current smoothed value. We define the smoothing and trend estimation at iteration i as:

$$s^{(i)} = \alpha r^{(i)} + (1 - \alpha) (s^{(i-1)} + b^{(i-1)}), \quad (1)$$

$$b^{(i)} = \beta (s^{(i)} - s^{(i-1)}) + (1 - \beta) b^{(i-1)}, \quad (2)$$

where the raw input is represented by $r^{(i)}$, the smoothed value is $s^{(i)}$, and $b^{(i)}$ is the best estimate of the trend. The parameters




Figure 6: Predictive steering methodology that maximizes the number of steers within a time period T , taking into account the estimated initial steering time $\bar{\epsilon}$, image processing and prediction overhead Δ , and subsequent steering overheads $\epsilon_{\{1..n-1\}}$.

$\alpha, \beta \in [0, 1]$ are the smoothing parameters that control the weight given to current and past values. In general, α and β values can be chosen with empirical guidance⁴, where larger values assign stronger weights to recent inputs and are suitable for faster motions. Given the smoothed value $s^{(i)}$ and estimated trend $b^{(i)}$, the *next* prediction is defined as $s^{(i+1)} = s^{(i)} + b^{(i)}$ occurring at iteration $i + 1$.

Figure 6 and Algorithm 2 describe the detailed steps. Once an image frame is captured at time t_0 , the system predicts the target location (x_T, y_T) at the next frame. The predicted location is produced at time t_1 after image processing delay Δ . The system then utilizes linear interpolation to infer locations between the captured pixel coordinate (x_0, y_0) and predicted coordinate (x_T, y_T) . It consecutively steers laser beam to these locations, taking into account the steering delay. Specifically, a steering command sent at t_i steers the laser beam to the interpolated location expected at time $t_i + \epsilon_i$. Since ϵ_i is not known before the command at t_i completes, we steer to the *expected* location based on the average steering delay $\bar{\epsilon} \approx \epsilon_i$. Since Lasertag’s mirror moves only a small amount between subsequent steers, $\bar{\epsilon}$ is a suitable approximation for ϵ (see §5.3).

4 Prototype Implementation

The Lasertag prototype consists of the core unit, containing the imaging and laser-steering subsystems, and a separate marker. The marker is a 7 cm diameter circle made of retroreflective tape [107]. We cut a 2 cm diameter hole in the center to place a photodiode (PD) [50] for performance evaluation.

Imaging Subsystem. The imaging subsystem images the scene to identify the marker’s location. Two 520 nm LEDs [93] are placed next to the system’s fisheye lens to flood the scene with green light. The green light reflects off the retroreflective marker and passes through the fisheye lens. The light then traverses the optical beamsplitter [130] and a 633 nm notch filter [105], which removes any stray reflections from the laser steering subsystem. Next, light passes through a 520 nm bandpass filter [125] with a 55 nm pass region bandwidth. The monochrome light is focused using an $f = 16$ mm focal length lens [106], before finally arriving at a 5 MP image sensor [102]. The processing pipeline is built upon the OpenMV H7 Plus

⁴Other prediction algorithms (see §5.3) can eliminate parameter tuning.




Figure 7: The Lasertag prototype and resulting beam profile measured at different distances.

platform. We perform adaptive windowing by sending IOCTL commands directly to the image sensor.

Laser-Steering Subsystem. The laser steering subsystem takes the marker location from the imaging subsystem and feeds the appropriate steering voltages to a digital-input MEMS driver [91] operating with an SPI bus frequency of 25 MHz. We utilize a 638 nm laser diode [129] with the driving current capped to 120 mA, resulting in 50 mW of optical power. The outgoing laser light is focused using an $f = 2.75$ mm aspheric lens [128]. After the aspheric lens, the laser light reflects off a 2.4 mm MEMS mirror [88] with a steering bandwidth of 300 Hz. Finally, the laser light passes through an $f = 20$ mm triplet lens [133]. The laser light then reflects off the beamsplitter and through the fisheye lens [103], achieving a 3 dB half-angle (x, y) divergence of (0.6, 0.8) degrees.

5 Evaluation

We evaluate Lasertag’s tethering efficacy in the context of three practical applications, followed by micro-benchmarks and the impact of practical factors. A demonstration is available at [15].

5.1 Experimental Setup

Since application-specific performance (e.g., communication, power delivery) fundamentally relies on the ability to maintain a constant laser connection, we primarily focus on evaluating Lasertag’s *tethering efficacy*. We then translate these results to application specific performance with additional experiments.

Ground Truth. We assemble a separate optical unit to collect ground truth marker and laser beam positions (Figure 8a). We combine two 2.4 MP cameras [138] with a 50:50 beamsplitter and place a fisheye lens [104] on the beamsplitter’s output for imaging. Each camera has a 520 nm or 633 nm bandpass filter for isolating the marker or laser spot. We add two 520 nm LEDs to generate retroreflective light. Each camera streams images to a host computer, in parallel, at an average 315 FPS. Frames from both cameras are synced by minimizing the difference between frame timestamps, and aligned by applying a homographic transform. We perform binarization, blob detection, and ellipse contouring to localize the laser spot and marker in each frame. The marker velocity is then estimated by taking the Euclidean distance between the marker’s center in adjacent frames and dividing by the frame interval. Notably, we mea-




Figure 8: Experimental setup for evaluating Lasertag in the context of three different practical scenarios.

sure *pixel velocity* instead of physical velocity since Lasertag’s performance is only dependent on pixel changes. Since pixel velocities scale with distance (e.g., a fast-moving object close to Lasertag has a large pixel velocity, and a slower pixel velocity when it is farther away), we can generalize Lasertag’s performance to any distance. We also add an accelerometer [99] to the target to complement the velocity data.

Evaluation Metrics. We measure tethering efficacy via: (1) the physical offset between the center of the laser beam and marker; (2) the normalized laser power received by a PD colocated with the marker.

We compute the first metric from ground-truth videos. For each ground truth frame, we compute the pixel offset as the Euclidean distance between the laser and marker blobs in the image. To convert pixel offset to physical offset, we approximate the focal length f of the optical ground-truth camera using the MATLAB and a conventional checkerboard calibration pattern. Assuming minimal lens distortion and a fixed marker depth d relative to the image plane, the physical offset o (in meters) corresponding to a pixel offset p (in pixels) is $o = p * d * 1/f$ [79], where f is experimentally measured to be 322.5 pixels. Notably, this equation can also be used to relate pixel velocity to physical velocity, where o is in m/s and p is in px/s. Since transient marker occlusions and image processing noise impact our measurements, we extract the most *likely* offset by applying a rolling median window of width 100 px/s over the velocity-offset pairs. We also ignore velocities higher than the 98th percentile (710 px/s) and lower than the 2nd percentile (13 px/s) due to lack of samples.

The second metric is measured by a PD placed at the center of the marker on the target. Throughout target motion, accelerometer data and PD voltage are sampled at 1 ksp/s by a micro-controller [5]. PD voltage and velocity data are then synced together to associate each PD reading with a velocity. We then perform a series of signal processing techniques to remove the confounding factors of depth, angular response, laser power, and occlusion – all of which detrimentally affect the PD’s response but do not indicate poor tethering performance. First, we observe that signal fluctuations due to target motion occur with a frequency between 3 Hz and 500 Hz. Thus we perform a 0.5 Hz rolling average of the signal’s 99th percentile, removing signal spikes caused by PD/ADC noise. We then rescale the readings between the 99th percentile rolling average and the PD’s noise floor (83 mV), partially removing the impact of the PD’s angular response, experiment

depth, and laser power. To account for temporary PD occlusion, we apply a rolling median window of width 100 px/s, and ignore velocities above 710 px/s and below 13 px/s.

Evaluation Procedure. We evaluate Lasertag in an enclosed indoor space. To understand the contribution of predictive steering, we also evaluate Lasertag without predictive steering as our baseline. First, we place the Lasertag core unit and ground truth measurement device on a 0.9 m height table (Figure 8a). Second, we place the marker and sensing components on the target (Figure 8). Third, we turn on a 1100 lx floor lamp (measured 1 m away) to provide ambient light and guidance for the drone’s optical positioning. Finally, we place the target at a fixed depth⁵ of either 0.8 m, 2.4 m, or 4 m from the Lasertag core unit, and record 3-8 minute trials for both steering methods at each depth. We synthesize over 1.6 million velocity measurements with associated PD readings. Laser safety requirements are maintained for all experiments, and since the laser’s reflections are mainly retroreflective, authors were only exposed to diffuse, eye-safe reflections.

Translating Tethering Efficacy. We translate tethering results to communication performance by mapping the marker PD’s voltage to the BER of a 1 Gbps stationary link between Lasertag and an avalanche photodiode (APD) [124]. First, we generate a random 1024-bit OOK, non-return-to-zero packet with an arbitrary waveform generator [62] connected to the Lasertag laser diode. Second, we colocate the APD and marker PD using a 50:50 beamsplitter, ensuring both sensors receive an identical copy of the modulated laser light. Third, we simultaneously capture the APD’s received signals using an oscilloscope [63] and the corresponding voltage of the marker PD. For each packet, we decode the signal in MATLAB, compute the BER, and measure the marker PD’s voltage. We then vary the link’s SNR by programmatically decreasing laser power. To extrapolate to all possible voltages, we fit all recorded samples with a sigmoid function (R^2 of 0.992). We then apply this function to our experimental PD values to obtain a BER over time. This methodology avoids placing the wired APD on our experimental targets which would restrict their motion. Translating tethering results to power-delivery performance is straightforward, as harvested power is directly proportional to connection time. We avoid reporting charging efficacy as it highly depends on the specific laser and solar cell.

⁵The max distance is greater than this depth, since path length increases at extreme steering angles. This is accounted for with our adaptive normalization.




Figure 9: (a-b) Distribution and velocity trends of received normalized power and physical offset across all scenarios. (c) Velocity distribution from individual scenarios. (d-f) Velocity breakdowns for each scenario.

5.2 Tethering Efficacy

To examine tethering efficacy, we consider object motions in three applications well-primed for the benefits of laser light: (1) ground-to-drone laser communication, (2) laser communication for VR, and (3) mobile power delivery using lasers. Across scenarios, our targets underwent realistic motions including unpredictable movements, jitter, velocity swings, and rotations.

Performance Overview. Across all applications, predictive steering consistently aligns the laser beam center with the PD on the target, delivering a median normalized laser power of 0.97 on average, compared to only 0.85 without predictive steering. Similarly, the median offset seen with predictive steering is 1.03 cm on average, compared to 1.17 cm for without predictive steering. To maintain the laser tether in the face of steering errors, the laser beam divergence could be expanded to increase the beam spot size and compensate for steering offsets, at the expense of energy efficiency. The required half-angle beam divergence to maintain the link in the face of a 1.17 cm offset is $\approx 0.67^\circ$ for a marker depth of 1 m. This low beam divergence corresponds to a relatively well-collimated beam, and thus Lasertag can efficiently deliver the majority of outputted laser energy to the marker’s receiver. Histograms of received normalized power and offsets across all scenarios are shown in Figure 9a, demonstrating predictive steering’s performance gain. We attribute this gain to the boosted steering rate of our predictive steering algorithm, achieving average rates up to 151 Hz from the baseline 36 Hz. As shown in Figure 9b, while predictive steering maintains relatively constant performance across velocities, baseline steering performance considerably declines. At the maximum observed marker speed of 710 px/s, corresponding to a physical speed of 8.8 m/s at a depth of 4 m, predictive steering achieves, on average, a 32% higher median power and 26% lower median offset than baseline steering.

In the context of power delivery, these results indicate that

Lasertag is able to deliver 97% of its laser power to mobile device, compared to 85% with baseline steering. Translating these results to communication performance, we observe a median BER of 1.15×10^{-8} ($p_{25} = 8.19 \times 10^{-9}$, $p_{75} = 7.97 \times 10^{-5}$) for predictive steering and 6.25×10^{-7} ($p_{25} = 9.26 \times 10^{-9}$, $p_{75} = 3.97 \times 10^{-4}$) for baseline steering. Notably, this corresponds to a 54-fold improvement in BER utilizing predictive steering. Including periods when the link was fully occluded, we observe *mean* BERs of 2.62×10^{-2} and 5.03×10^{-2} for predictive and baseline steering, respectively. To improve these results, the laser power can be increased to improve the overall SNR (thereby decreasing BER), or lower throughputs (e.g., 100s of Mbps) can be considered to relax the SNR requirements.

Diving into each application, we observe that Lasertag’s performance is highly dependent on not only the target’s velocity, but also its trajectory. We identify two broad classes of object motion: (1) continuous linear movement; (2) short periods of linear or angular movement, which register as bursts of small, seemingly random displacements (later referred to as micro-movements). Importantly, both continuous and micro-movements can occur at all velocities. For instance, although the drone predominately flies in a continuous, linear fashion at a variety of speeds, VR games generate a combination of slow and rapid angular head movements, generating micro-movements of varying velocities. The presence of continuous, linear motion vs. micro-movements affects the received power and offset achieved by both steering methods.

Lastly, we observe that the distribution of velocities across applications is heavily positively skewed (Figure 9c). Since there are fewer samples recorded at higher velocities, we observe a larger variance in received power and offset at higher velocities. Despite this, predictive steering maintains a tighter confidence interval around the median than baseline steering. Additionally, a slightly lower PD reading and higher

offset exists at the lowest pixel velocities which we attribute to occlusion not accounted for by our adaptive normalization. Below we break down our results by application and examine the impact of specific motions on Lasertag's tethering efficacy.

Drone Communication. We consider maintaining a communication link with a moving aerial drone for our first application scenario. We track, in real-time, a DJI Mini 2 drone flying with a maximum acceleration of 29 m/s^2 in the z-direction and 19 m/s^2 in the x-y plane. At each starting depth, the drone flew left-to-right in repeated sweeps across the room, each measuring roughly 5 m horizontally. With the exception of take-off and landing, the drone maintained a relatively constant vertical position. As shown in Figure 9c, a significant number of samples at low velocities were collected, corresponding to instances in which the drone was changing direction, such that the gradient of pixel movement was close to zero. The second peak and tail of the distribution correspond to samples collected while the drone cruised across the room.

On average across all measured velocities, predictive steering delivers 27% higher median-normalized laser power than baseline steering and maintains a 22% smaller median offset. As the marker velocity increases beyond 230 px/s, predictive steering maintains median normalized powers above 0.86 and median offsets less than 1.11 cm. In contrast, baseline steering performance significantly declines, such that at the maximum recorded velocity of 710 px/s it achieves a median normalized power of 0.71 and median 1.61 cm offset. Translating these results to communication performance, we observe a median BER of 9.26×10^{-9} for predictive steering and 2.73×10^{-6} for baseline steering. Notably, a large portion of the drone's motion consists of continuous, linear motion, which is ideal for double exponential prediction. As such, we observe increased tethering performance with predictive steering enabled.

Laser VR. For our second scenario, we consider laser communication to a VR headset for high-quality content streaming. We track a Pico4 VR headset while the participant plays a game, moving left/right up to 5 m, up/down up to 1.8 m, and back/forth up to 5 m. Throughout all trials, the VR headset experienced a maximum acceleration of 37 m/s^2 in the z-direction and 15 m/s^2 in the x-y plane.

As shown in Figure 9e, predictive steering and baseline steering achieve median normalized laser powers of 0.99 and 0.98 and offsets of 0.81 cm and 0.89 cm, respectively, with relatively constant performance across velocities. The noticeably smaller performance gain of predictive steering is due to the prominence of micro-movements throughout VR gameplay. While such movements are typically not amenable to predictive steering, here they correspond to small (albeit fast) *physical* displacements of the marker such that the PD remains within the laser beam's diameter at all times. Steering without prediction similarly benefits from this motion pattern. Translating these results to communication performance, we observe a median BER of 1.11×10^{-8} for predictive steering and 1.18×10^{-8} for baseline steering. For high-resolution VR requiring multiple

Gbps of throughput, the laser power can be increased to improve the overall SNR and decrease the resulting BER.

Mobile Power Delivery. For our final scenario, we consider the potential of continuously delivering power to mobile devices. We track the back of a smartphone while a participant performs various actions, including (1) talking on the phone while walking, (2) pausing to answer a video call, and (3) bending over while on the phone. Throughout trials, the participant moved left/right up to 5 m, up/down up to 1.8 m, and back/forth up to 5 m. The phone experienced a maximum acceleration of 49 m/s^2 in the z-direction and 27 m/s^2 in the x-y plane.

As shown in Figure 9f, predictive steering on average delivers 19% higher median normalized laser power. Predictive steering consistently achieves median normalized laser powers above 0.90. In the context of power delivery, this translates to a mobile device receiving over 90% of laser energy from the core unit. The median offset achieved by baseline steering on average is 1.33 cm, compared to 1.25 cm with predictive steering. As in the case of drone motion, higher velocities see increased performance differences between predictive and baseline steering. Additionally, during slower motions, the user occasionally occluded the PD, causing a consistent decline in power at velocities below 300 px/s that negatively skewed the median.

5.3 Tethering Micro-benchmarks

Motion Prediction. To assess the accuracy of motion prediction, we compare the predicted marker pixel locations with ground truth values. We first downsample the ground truth values (recorded at 315 FPS) to the average frame rate of Lasertag's image sensor (36 FPS). We then compare our prediction algorithm to two reference points: (1) *baseline without prediction*, which treats the target's current location as its next location; (2) *optimal case*, which uses a brute-force search to find the optimal hyperparameters α, β of the double exponential filter for each application. Although the optimal case is not practical for real-time applications due to the time-consuming brute-force search, we include it as an upper bound for comparison. We quantify prediction error as the Euclidean distance between the predicted and ground truth marker pixel locations and report the mean/STDV of prediction error.

As shown in Figure 10a, compared to the baseline, our predictive steering algorithm reduces prediction error by roughly half in all applications, achieving an average error of 1.67 px, corresponding to a 1.04 cm error at a distance of 2 m. The difference between our prediction algorithm to the optimal case is, on average, only 0.3 px (1.8 mm at 2 m). Across prediction methods, we observe that applications with higher marker velocities (i.e., drone motion) see higher mean prediction errors. This relationship arises because higher marker velocities correspond to larger displacements of the marker between consecutive frames. Inaccurate predicted locations are then increasingly far from the actual location. Tuning of the double exponential parameters or more sophisticated motion prediction models (§6) could improve accuracy at higher speeds.

Angular Range. Next we measure the angular range of the




Figure 10: Lasertag micro-benchmarks and practical considerations including (a) prediction accuracy, (b) angular range of the steering and imaging subsystems, (c) the timing breakdown, and (d) impact of background noise.

laser subsystem and the imaging subsystem. To quantify the imaging range, we place a power meter [131] at the center of the retroreflective marker to measure optical power. We then rotate the marker in 10° increments when 18 cm away from Lasertag and measure the power at each angle. To measure the laser subsystem’s range, we measure the *maximum* received power at each angle, indicating how efficient the optical circuit is at delivering laser light to any angle. In both scenarios, we measure the received power relative to the maximum power.

As shown in Figure 10b, the laser subsystem is capable of delivering over 56% of its laser power up to $\pm 95^\circ$, covering 10° more than a full hemisphere before dropping below its -3 dB point. Notably, the steering is slightly asymmetrical around the 0° mark. We hypothesize this is due to reduced MEMS mirror area at extreme positive angles, which lowers reflection efficiency [90]. The imaging subsystem delivers $\geq 66\%$ of its power to the marker center over a $\pm 81^\circ$ range. This is because at extreme angles, the circular retroreflective marker becomes distorted, resulting in an incorrect center measurement.

System Speed. Finally we evaluate the delay of each step in Lasertag and plot the timing breakdown in Figure 10c. Overall, image capture takes the most time (27.8 ms), limited by the frame rate (36 FPS). Lasertag runs other steps in parallel with image capture, including target tracking (7.9 ms), motion prediction and interpolation (≈ 0.1 ms), and consecutive steering (2.1 ms each on average). On average, these timing constraints support up to four predictive steers before the next frame is requested. For a tag moving with an angular speed of $1^\circ/\text{s}$ (roughly corresponding to $10 \text{ px}/\text{s}$), the MEMS mirror takes approximately 0.6 ms to rotate to its updated physical deflection angle. The angular response is nonlinear, maxing out at ≈ 6 ms for angles above 30° . Notably, all these durations are extremely variable and dependent on hardware characteristics, image sensor configuration, and marker movement. Consequently, we observe a variable frame rate from 238 FPS (when the marker is stationary) to 14 FPS (when the marker is causing frequent ROI window updates). As a result, the number of predictive steers varies from an observed minimum of one to a maximum of eighteen. Faster MCUs and a streamlined MEMS communication protocol would decrease the overall time spent in image processing and communicating with the mirror, thereby enabling more steering between frames.

5.4 Practical Considerations

Finally, we investigate practical aspects of Lasertag.

Power Consumption. Using a Monsoon FTA22D power meter, we measure the power consumption of Lasertag (Table 1). Overall, the power consumption of the MCU, image sensor, and MEMS mirror is less than 1.5 W, 70% smaller than the *minimum* power consumption of common computer vision processors, e.g., the 5 W Jetson Nano. Notably, our chosen laser consumes 353 mW of input power to generate 50 mW of optical power. Higher efficiency/lower-power diodes can fit seamlessly within the system to lower the power consumption.

Noisy Environments. To examine Lasertag’s robustness against confounding scene objects, we directly measure the perceived brightness of the retroreflective marker (cut to three diameters) and common scene objects (e.g. reflective screens, glass, whiteboards). Figure 10d plots the mean brightness of the marker as function of distance. In general, larger markers retroreflect more light than smaller markers, resulting in brighter blobs on the image sensor.⁶ In general, fairly low-power LEDs have comparable brightness to the retroreflective markers, supporting our claim that active markers would improve Lasertag’s overall power efficiency at the cost of system complexity. Additionally, these results validate the importance of the marker detection algorithm in §3.2 which enables Lasertag to operate with bright light sources in the scene.

Ambient Light. Similar to supported imaging range, ambient light rejection is tightly coupled with the chosen optical filters. To reject more ambient light, a narrow bandpass filter can be utilized to only pass a small subset of wavelengths. In our implementation, we opted for a wider bandwidth (55 nm) filter to match the broad emission spectrum of the 520 nm LEDs. Despite the wide bandwidth, it still required over 200,000 lx of light to saturate the image sensor using a white LED.

6 Discussion

Additional Tracking Modes. Our framework is compatible with any image-based detection technique and is thus application-agnostic. The tracking system may be configured to support two additional modes: (1) tracking active markers, and (2) markerless object tracking. An active marker can be

⁶We observe a decrease in blob brightness within 1 m of Lasertag due to the majority of retroreflective light returning to their source instead of overlapping with the colocated fisheye lens.

Table 1: Average Lasertag power consumption breakdown.

Imaging	Power (mW)	Steering	Power (mW)
MCU	596	50mW LD	353
Image sensor	692	MEMS mirror	159
2x LEDs	305		
Subtotal	1593 mW	Subtotal	512 mW

realized by adding an active light emitter (e.g., IR LED) to the target, potentially enabling alignment feedback or uplink communication (e.g., using frequency modulation). This would also increase sensing ranges and robustness to background objects or targeted attacks (e.g., use of "spoofing" markers to fool Lasertag). The second mode requires no marker on the object, and relies fully on either standard feature detection algorithms or machine learning models to identify the target. Prior work has examined the tracking of people [10], traffic [24], and arbitrary objects [122], which can be plugged into the Lasertag framework. We leave these explorations to future work.

Motion Prediction. More sophisticated prediction models can predict the target's next position with higher accuracy. Recent work has proposed learning models that can be trained to learn complex patterns and relationships between an object's motion and its surrounding environment [13, 38, 117]. To test the potential benefits of deep learning prediction for Lasertag, we apply the N-euro Predictor [117] to the data collected during our experiments without further fine-tuning. This model reduced the prediction error from an average of 1.67 px to 0.8 px, a 52% decrease. This performance boost requires more powerful hardware for neural network inference, which Lasertag can support depending on the application requirements.

Path Blockage. Lasertag can support non-line of sight applications by leveraging its image sensor to find reflections in the scene and peer around obstructions. For example, imagine a scenario where the Lasertag core unit and a target are separated by a wall, but to one end of the wall is a mirror that is visible by both. The core unit will see the target's reflection in the mirror, and can steer the laser beam to the reflection. Because of path symmetry between the image sensor and laser, the beam will reflect off the mirror and onto the target.

7 Related Work

Colocating Tracking/Positioning with Laser Steering. Prior work has studied the integration of object acquisition, tracking, and laser steering to maintain connection of a laser beam with a target object. To the best of our knowledge, no existing systems integrate target-agnostic computer vision based tracking with laser steering. Existing systems often equip the target with retroreflectors [14, 59, 77, 83, 86, 87], photodiodes [87], or active emitters like LEDs [83, 87, 95, 134] or laser beacons [12, 20, 61, 119, 123, 139, 144]. Lasertag imposes no restrictions on the target's design, as it can be reconfigured to support tracking and steering to any object visible to the embedded camera. [4, 126] require a multi-camera system to determine the mirror tilt needed to steer to a 3D po-

sition. These systems must be carefully calibrated by hard-coding the camera system's parameters [4, 126]. Lasertag requires no *a priori* knowledge of camera characteristics, instead automatically determining the correct steering parameters during the mapping stage. [59, 83, 86, 87] perform iterative scans to track and steer to a passive mobile target. In contrast, Lasertag continuously steers directly to the target and thus does not suffer from tether downtime resulting from scanning. While [59, 83, 86, 87] perform several demonstrations of their system and describe its design, quantified performance results are not provided. [1, 2, 14, 39, 51, 64] rely on uplink feedback between the transmitter and receiver, either to enable initial beam alignment [1, 64, 141], continuously communicate link quality [39, 51] or receive GPS locations [2, 14]. We employ a purely optical approach that does not require uplink feedback.

Laser Steering. Laser beam steering mechanisms fall into two categories: non-mechanical and mechanical. Mechanical beam steering mechanisms involving rotating mirrors have large steering ranges but tend to be bulky, sensitive to mechanical stress, and limited in steering resolution [60, 142]. Non-mechanical beam steering mechanisms include acousto-optic deflectors [121, 136], optical gratings [11, 41, 67, 80, 100], tunable liquid lenses [22, 101, 150], phased arrays [29, 30, 35], and MEMS mirrors [33, 47, 85, 146]. Overall, MEMS mirror are the most appealing due to their small size, low cost, and robustness. Other solutions lead to significant loss of optical power during steering [121, 136], suffer from limited FoVs and switching speeds [11, 22, 41, 67, 80, 100, 101, 150], or have limited commercial availability [29, 30, 35].

3D Object Positioning and Tracking. [94] and [37] explore the use of feature matching combined with reprojection to perform 3D tracking. A large body of work explores light-based positioning and tracking techniques that do not employ cameras [42, 65, 76], including those requiring [46, 70, 110] active emitters or photodiodes [31, 52, 145] on the target. Other 3D object localization methodologies include ultrasonic [21], LiDAR [49], MEMS [73, 87], IMU [36, 72], and RF [3, 8, 57, 68, 120, 148] sensing. As discussed in §2, these methodologies are nontrivial to integrate with laser steering devices.

8 Conclusion

We present Lasertag, a reconfigurable framework for laser tethering with highly mobile targets. Extensive evaluation shows that Lasertag can maintain laser alignment with rapidly moving targets equipped with retroreflective markers. Lasertag paves the way for a myriad of laser-based mobile applications, including laser communication and wireless power delivery.

9 Acknowledgements

We sincerely thank our reviewers and shepherd for their insightful feedback. This work is supported in part by the National Science Foundation (GRFP-1840344, CNS-1552924). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the funding agencies or others.

References

- [1] Mojtaba Mansour Abadi, Mitchell A Cox, Rakan E Alsaigh, Shaun Viola, Andrew Forbes, and Martin PJ Lavery. A space division multiplexed free-space-optical communication system that can auto-locate and fully self align with a remote transceiver. *Scientific Reports*, 9(1):1–8, 2019.
- [2] Mouhammad Al Akkoui, Robert Huck, and James Sluss. High-speed communications enabling real-time video for battlefield commanders using tracked FSO. *Proc. of SPIE - The International Society for Optical Engineering*, 2007.
- [3] Marco Altini, Davide Brunelli, Elisabetta Farella, and Luca Benini. Bluetooth indoor localization with multiple neural networks. In *IEEE 5th International Symposium on Wireless Pervasive Computing*, pages 295–300, 2010.
- [4] Nicolas Andreff, Brahim Tamadazte, Sounkalo Dembélé, and Zill E Hussnain. Preliminary variation on multiview geometry for vision-guided laser surgery. In *Workshop on Multi-View Geometry in Robotics*, pages 1–10, 2013.
- [5] Arduino. Arduino MKR1000 WiFi. <https://store-usa.arduino.cc/products/arduino-mkr1000-wifi>, 2023.
- [6] Charles K Asawa. Boresighting system for infrared optical receiver and transmitter, May 2 1978. US Patent 4,087,689.
- [7] Ronald Azuma and Gary Bishop. Improving static and dynamic registration in an optical see-through HMD. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 197–204, 1994.
- [8] R. Bajaj, S.L. Ranaweera, and D.P. Agrawal. Gps: location-tracking technology. *Computer*, 35(4):92–94, 2002.
- [9] Sunil L Bangare, Amruta Dubal, Pallavi S Bangare, and Suhas Patil. Reviewing otsu’s method for image thresholding. *International Journal of Applied Engineering Research*, 10(9):21777–21783, 2015.
- [10] Antonio Brunetti, Domenico Buongiorno, Gianpaolo Francesco Trotta, and Vitoantonio Bevilacqua. Computer vision and deep learning techniques for pedestrian detection and tracking: A survey. *Neurocomputing*, 300:17–33, 2018.
- [11] J Buck, S Serati, R Serati, H Masterson, M Escuti, J Kim, and M Miskiewicz. Polarization gratings for non-mechanical beam steering applications. In *Acquisition, Tracking, Pointing, and Laser Systems Technologies XXVI*, volume 8395, page 83950F, 2012.
- [12] Jamie W Burnside, Stephen D Conrad, Allen D Pillsbury, and Catherine E DeVoe. Design of an inertially stabilized telescope for the Ilcd. In *Free-Space Laser Communication Technologies XXIII*, volume 7923, pages 133–140, 2011.
- [13] Judith Butepage, Michael J Black, Danica Kragic, and Hedvig Kjellstrom. Deep representation learning for human motion prediction and classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6158–6166, 2017.
- [14] Alberto Carrasco-Casado, Ricardo Vergaz, José M Sánchez-Pena, Eva Otón, Morten A Geday, and José M Otén. Low-impact air-to-ground free-space optical communication system design and first results. In *2011 International Conference on Space Optical Systems and Applications (ICSOS)*, pages 109–112, 2011.
- [15] Charles J Carver, Hadleigh Schwartz, Qijia Shao, Nicholas Shade, Joseph P Lazzaro, Xiaoxin Wang, Jifeng Liu, Eric R Fossum, and Xia Zhou. Catch me if you can: Demonstrating laser tethering with highly mobile targets. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, pages 1–3, 2023.
- [16] Charles J Carver, Qijia Shao, Samuel Lensgraf, Amy Sniffen, Maxine Perroni-Scharf, Hunter Gallant, Alberto Quattrini Li, and Xia Zhou. Sunflower: Locating underwater robots from the air. In *Proc. of MobiSys*, page 14–27, 2022.
- [17] Charles J Carver, Zhao Tian, Qijia Shao, Hongyong Zhang, Kofi M Odame, Alberto Quattrini Li, and Xia Zhou. Air-water communication and sensing with light. In *2022 14th International Conference on COMMunication Systems & NETWORKS (COMSNETS)*, pages 371–374, 2022.
- [18] Charles J Carver, Zhao Tian, Hongyong Zhang, Kofi M. Odame, Alberto Quattrini Li, and Xia Zhou. AmphiLight: Direct air-water communication with laser light. In *Proc. of NSDI*, pages 373–388, 2020.
- [19] Justin Chan, Ananditha Raghunath, Kelly E. Michaelsen, and Shyamnath Gollakota. Testing a drop of liquid using smartphone lidar. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 6(1), 2022.
- [20] C Chen, A Grier, M Malfa, E Booen, H Harding, C Xia, M Hunwardsen, J Demers, K Kudinov, G Mak, B Smith, A Sahasrabudhe, F Patawaran, T Wang, A Wang, C Zhao, D Leang, J Gin, M Lewis, B Zhang, D Nguyen,

- D Jandrain, F Haque, and K Quirk. Demonstration of a bidirectional coherent air-to-ground optical link. In *Free-Space Laser Communication and Atmospheric Propagation*, volume 10524, page 105240G, 2018.
- [21] Moi Tin Chew, Fakhru Alam, Mathew Legg, and Gourab Sen Gupta. Accurate ultrasound indoor localization using spring-relaxation technique. *Electronics*, 10(11):1290, 2021.
- [22] Hyun Choi and Wan-Chin Kim. Design of mechaless LiDAR optical system with large FOV using liquid lens and fisheye lens. In *ASME-JSME 2018 Joint International Conference on Information Storage and Processing Systems and Micromechatronics for Information and Precision Equipment*, page V001T10A001, 2018.
- [23] Hyun Choi, No-Cheol Park, and Wan-Chin Kim. Optical system design for light detection and ranging with ultra-wide field-of-view using liquid lenses. *Microsystem Technologies*, 26(1):121–131, 2020.
- [24] Benjamin Coifman, David Beymer, Philip McLauchlan, and Jitendra Malik. A real-time computer vision system for vehicle tracking and traffic surveillance. *Transportation Research Part C: Emerging Technologies*, 6(4):271–288, 1998.
- [25] Max Curran, Md Shaifur Rahman, Himanshu Gupta, Kai Zheng, Jon Longtin, Samir R. Das, and Thanvir Mohamed. FSONet: A wireless backhaul for multi-gigabit picocells using steerable free space optics. In *Proc. of MobiCom*, page 154–166, 2017.
- [26] Alan J. Danker and Azriel Rosenfeld. Blob detection by relaxation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-3(1):79–92, 1981.
- [27] Carlo De Santi, Matteo Meneghini, Alessandro Caria, Ezgi Dogmus, Malek Zegaoui, Farid Medjdoub, Boris Kalinic, Tiziana Cesca, Gaudenzio Meneghesso, and Enrico Zanoni. Gan-based laser wireless power transfer system. *Materials*, 11(1):153, 2018.
- [28] DJI. DJI Mini 3 Specs. <https://www.dji.com/mini-3/specs>, 2023.
- [29] Jonathan Doylend, Martijn R Heck, Jock Bovington, Jon Peters, Larry Coldren, and John Bowers. Free-space beam steering in two dimensions using a silicon optical phased array. In *Optical Fiber Communication Conference*, pages OM2J–1, 2012.
- [30] Jonathan K. Doylend, M.J.R. Heck, Jock T. Bovington, Jonathan D Peters, L.A. Coldren, and J.E. Bowers. Two-dimensional free-space beam steering with an optical phased array on silicon-on-insulator. *Optics Express*, 19(22):21595–21604, 2011.
- [31] Minzhen Du. *Assessment of a Low Cost IR Laser Local Tracking Solution for Robotic Operations*. PhD thesis, Virginia Tech, 2021.
- [32] Josef Eisenring, Klaus W Hildebrand, and Jakob Tanner. Coaxial transmitting and receiving optics for an electro-optic range finder, August 28 1979. US Patent 4,165,936.
- [33] Yusuf Said Eroglu, Ismail Guvenc, Alphan Sahin, Nezh Pala, and Murat Yuksel. Diversity combining and piezoelectric beam steering for multi-element VLC networks. In *Proceedings of the 3rd Workshop on Visible Light Communication Systems*, pages 25–30, 2016.
- [34] Zhufeng Fan, Jinyu Zhan, and Wei Jiang. Detecting deepfake videos by visual-audio synchronism: work-in-progress. In *Proceedings of the 2021 International Conference on Embedded Software*, pages 31–32, 2021.
- [35] Reza Fatemi, Aroutin Khachaturian, and Ali Hajimiri. A low power PWM optical phased array transmitter with 16° field-of-view and 0.8° beamwidth. In *2018 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, pages 28–31, 2018.
- [36] Ling feng Shi, Yu-Le Zhao, Gong xu Liu, Sen Chen, Yue Wang, and Yi-Fan Shi. A robust pedestrian dead reckoning system using low-cost magnetic and inertial sensors. *IEEE Transactions on Instrumentation and Measurement*, 68:2996–3003, 2019.
- [37] Dmitry Gorodnichy, S Malik, and Gerhard Roth. Affordable 3D face tracking using projective vision. In *Proc. of Int. Conf. on Vision Interface*, pages 383–390, 2002.
- [38] Mahir Gulzar, Yar Muhammad, and Naveed Muhammad. A survey on motion prediction of pedestrians and vehicles for autonomous driving. *IEEE Access*, 9:137957–137969, 2021.
- [39] Himanshu Gupta, Max Curran, Jon Longtin, Torin Rockwell, Kai Zheng, and Mallesham Dasari. Cyclops: An FSO-Based Wireless Link for VR Headsets. In *Proc. of SIGCOMM*, page 601–614, 2022.
- [40] Markus Haase and Helmut Schäfer. Upconverting nanoparticles. *Angewandte Chemie International Edition*, 50(26):5808–5829, 2011.
- [41] Michael Hall, Qing Chao, Byron Taylor, and Xinqiao Liu. Non-mechanical beam steering for depth sensing, February 25 2020. US Patent 10,574,973.
- [42] Anum Hameed and Hafiza Anisa Ahmed. Survey on indoor positioning applications based on different technologies. In *2018 12th International Conference on Mathematics, Actuarial Science, Computer Science and Statistics (MACS)*, pages 1–5, 2018.

- [43] Yuze He, Li Ma, Zhehao Jiang, Yi Tang, and Guoliang Xing. Vi-eye: Semantic-based 3D point cloud registration for infrastructure-assisted autonomous driving. In *Proc. of MobiCom*, page 573–586, 2021.
- [44] Niels Henze, Sven Mayer, Huy Viet Le, and Valentin Schwind. Improving software-reduced touchscreen latency. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*, 2017.
- [45] S Hinz. Fast and subpixel precise blob detection and attribution. In *IEEE International Conference on Image Processing*, volume 3, pages III–457, 2005.
- [46] Andrew Hogue, Matt Robinson, MR Jenkin, and Robert S Allison. A vision-based head tracking system for fully immersive displays. In *Proceedings of the workshop on Virtual environments*, pages 179–187, 2003.
- [47] Sven TS Holmstrom, Utku Baran, and Hakan Urey. MEMS laser scanners: a review. *Journal of Microelectromechanical Systems*, 23(2):259–275, 2014.
- [48] Jinhui Huang, Chunlin Li, and Jie Yu. Resource prediction based on double exponential smoothing in cloud computing. In *2012 2nd International Conference on Consumer Electronics, Communications and Networks*, pages 2056–2060, 2012.
- [49] Rui Huang, Wanyue Zhang, Abhijit Kundu, Caroline Rebecca Pantofaru, David Alexander Ross, Thomas Funkhouser, and Alireza Fathi. An LSTM approach to temporal 3d object detection in lidar point clouds. In *ECCV*, pages 266–282, 2020.
- [50] Texas Instruments. OPT101. <https://www.ti.com/lit/ds/symlink/opt101.pdf>, 2023.
- [51] Femi Ishola and Mengu Cho. Experimental study on photodiode array sensor aided mems fine steering mirror control for laser communication platforms. *IEEE Access*, 9:100197–100207, 2021.
- [52] Shahidul Islam, Bogdan Ionescu, Cristian Gadea, and Dan Ionescu. Indoor positional tracking using dual-axis rotating laser sweeps. In *IEEE International Instrumentation and Measurement Technology Conference Proceedings*, pages 1–6, 2016.
- [53] David J Israel, Bernard L Edwards, Richard L Butler, John D Moores, Sabino Piazzolla, Nic du Toit, and Lena Braatz. Early results from NASA’s laser communications relay demonstration (LCRD) experiment program. In *Free-Space Laser Communications XXXV*, volume 12413, page 1241303, 2023.
- [54] Vikram Iyer, Elyas Bayati, Rajalakshmi Nandakumar, Arka Majumdar, and Shyamnath Gollakota. Charging a smartphone across a room using lasers. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 1(4), January 2018.
- [55] Johannes James, Vikram Iyer, Yogesh Chukewad, Shyamnath Gollakota, and Sawyer B. Fuller. Liftoff of a 190 mg laser-powered aerial vehicle: The lightest wireless robot to fly. In *Proc. of ICRA*, 2018.
- [56] Tim P Johnson and Jose Sasian. Image distortion, pupil coma, and relative illumination. *Appl. Opt.*, 59(22):G19–G23, 2020.
- [57] Kiran Joshi, Dinesh Bharadia, Manikanta Kotaru, and Sachin Katti. WiDeo: Fine-grained Device-free Motion Tracing using RF Backscatter. In *Proc. of NSDI*, page 189–204, 2015.
- [58] Jordin Kare and Thomas Nugent. Laser power beaming on a shoestring. *AIP Conference Proceedings*, 997:97–108, 04 2008.
- [59] Abhishek Kasturi, Veljko Milanovic, Bryan H Atwood, and James Yang. UAV-borne lidar with mems mirror-based scanning capability. In *Laser Radar Technology and Applications XXI*, volume 9832, pages 206–215, 2016.
- [60] Yagiz Kaymak, Roberto Rojas-Cessa, Jianghua Feng, Nirwan Ansari, MengChu Zhou, and Tairan Zhang. A survey on acquisition, tracking, and pointing mechanisms for mobile free-space optical communications. *IEEE Communications Surveys and Tutorials*, 20(2):1104–1123, 2018.
- [61] Kamugisha Kazaura, Kazunori Omae, Toshiji Suzuki, Mitsuji Matsumoto, Edward Mutaungwa, Tadaaki Murakami, Koichi Takahashi, Hideki Matsumoto, Kazuhiko Wakamori, and Yoshinori Arimoto. Performance evaluation of next generation free-space optical communication system. *IEICE Transactions on Electronics*, E90-C, 2007.
- [62] Keysight. M8190A. <https://www.keysight.com/us/en/product/M8190A/12-gsa-s-arbitrary-waveform-generator.html>, 2023.
- [63] Keysight. MSOS254A. <https://www.keysight.com/us/en/product/MSOS254A/high-definition-oscilloscope-2-5ghz-4-analog-16-digital-channels.html>, 2023.
- [64] Mahmudur Khan and Murat Yuksel. Autonomous alignment of free-space-optical links between UAVs. In *Proc. of HotWireless*, 2015.

- [65] Hiam M Khoury and Vineet R Kamat. Evaluation of position tracking technologies for user localization in indoor construction environments. *Automation in construction*, 18(4):444–457, 2009.
- [66] Dennis K. Killinger. Free space optics for laser communication through the air. *Optics & Photonics News*, 13:36–42, 2002.
- [67] AMJ Koonen, CW Oh, and E Tangdiongga. Reconfigurable free-space optical indoor network using multiple pencil beam steering. In *2014 OptoElectronics and Communication Conference and Australian Conference on Optical Fibre Technology*, pages 204–206, 2014.
- [68] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. Spotfi: Decimeter level localization using wifi. *Proc. of SIGCOMM*, 45(4):269–282, 2015.
- [69] A. Kotsopoulos, Pantazi A., and Antonakopoulos T. Control for high-speed archimedean spiral nanopositioning. In *2010 17th IEEE International Conference on Electronics, Circuits and Systems*, page 3, 2010.
- [70] Anil Kumar and Pinhas Ben-Tzvi. Spatial object tracking system based on linear optical sensor arrays. *IEEE Sensors Journal*, 16(22):7933–7940, 2016.
- [71] Nina Siu-Ngan Lam. Spatial interpolation methods: a review. *The American Cartographer*, 10(2):129–150, 1983.
- [72] Ariel Larey, Eliel Akin, and Itzik Klein. Multiple inertial measurement units—an empirical study. *IEEE Access*, PP:1–1, 04 2020.
- [73] Steven M LaValle, Anna Yershova, Max Katsev, and Michael Antonov. Head tracking for the Oculus Rift. In *Proc. of ICRA*, pages 187–194, 2014.
- [74] Joseph J. LaViola. Double Exponential Smoothing: An Alternative to Kalman Filter-Based Predictive Tracking. In *Proceedings of the workshop on Virtual environments*, page 199–206, 2003.
- [75] Xiaobao Lee and Chunhui Wang. Optical design for uniform scanning in MEMS-based 3D imaging LiDAR. *Appl. Opt.*, 54(9):2219–2223, 2015.
- [76] Angus Leigh, Joelle Pineau, Nicolas Olmedo, and Hong Zhang. Person tracking and following with 2D laser scanners. In *Proc. of ICRA*, pages 726–733, 2015.
- [77] Long Li, Runzhou Zhang, Zhe Zhao, Guodong Xie, Peicheng Liao, Kai Pang, Haoqian Song, Cong Liu, Yongxiong Ren, Guillaume Labroille, Pu Jian, Dmitry Starodubov, Brittany Lynn, Robert Bock, Moshe Tur, and Alan E. Willner. High-capacity free-space optical communications between a ground transmitter and a ground receiver via a uav using multiplexing of multiple orbital-angular-momentum beams. *Scientific Reports*, 7(1):17427, 2017.
- [78] Jiandong Liang, Chris Shaw, and Mark Green. On temporal-spatial realism in the virtual reality environment. In *Proc. of UIST*, page 19–25, 1991.
- [79] Kuan-Ying Lin, Yi-Hsing Tseng, and Kai-Wei Chiang. Interpretation and transformation of intrinsic camera parameters used in photogrammetry and computer vision. *Sensors*, 22(24), 2022.
- [80] Yu-Hua Lin, Milind Mahajan, Donald Taber, Bing Wen, and Bruce Winker. Compact 4 cm aperture transmissive liquid crystal optical phased array for free-space optical communications. In *Free-Space Laser Communications V*, volume 5892, page 58920C, 2005.
- [81] Wolfram MathWorld. Archimedes’ spiral. <https://mathworld.wolfram.com/ArchimedesSpiral.html>, 2021.
- [82] Joanna McKittrick and Lauren E Shea-Rohwer. Down conversion materials for solid-state lighting. *Journal of the American Ceramic Society*, 97(5):1327–1352, 2014.
- [83] V Milanović, N Siu, A Kasturi, M Radojičić, and Y Su. MEMSEye for optical 3D position and orientation measurement. In *MOEMS and Miniaturized Systems X*, volume 7930, page 79300U, 2011.
- [84] Veljko Milanovic, Kenneth Castelino, and Daniel T. McCormick. Highly adaptable MEMS-based display with wide projection angle. In *IEEE 20th International Conference on Micro Electro Mechanical Systems (MEMS)*, pages 143–146, 2007.
- [85] Veljko Milanovic, Kenneth Castelino, and Daniel T. McCormick. Highly adaptable MEMS-based display with wide projection angle. In *IEEE 20th International Conference on Micro Electro Mechanical Systems*, pages 143–146, 2007.
- [86] Veljko Milanović, Abhishek Kasturi, James Yang, and Frank Hu. A fast single-pixel laser imager for vr/ar headset tracking. In *MOEMS and Miniaturized Systems XVI*, volume 10116, page 101160E, 2017.
- [87] Veljko Milanovic and Wing Kin Lo. Fast and high-precision 3D tracking and position measurement with mems micromirrors. In *IEEE/LEOS International Conference on Optical MEMs and Nanophotonics*, pages 72–73, 2008.
- [88] MirrorcleTech. A5M24.3-2400AL. <https://www.mirrorcletech.com/pdf/DSI/>

- MirrorcleTech_Datasheet_A5M24.3-2400AL.pdf, 2023.
- [89] MirrorcleTech. AN002 - approximate linear models for mems mirrors. https://www.mirrorcletech.com/pdf/AN/AN002_-_Approximate_Linear_Models_for_MEMS_Mirrors.pdf, 2023.
- [90] MirrorcleTech. Angle of incidence and field distortions. https://www.mirrorcletech.com/pdf/AN/AN004_-_Angle_of_Incidence_and_Field_Distortions.pdf, 2023.
- [91] MirrorcleTech. DR-10-056-00. <https://www.mirrorcletech.com/wp/products/hardware/drivers/digital/>, 2023.
- [92] Steve P Monacos, Angel A Portillo, William Liu, James W Alexander, and Gerardo G Ortiz. A high frame rate CCD camera with region-of-interest capability. In *2001 IEEE Aerospace Conference Proceedings (Cat. No. 01TH8542)*, volume 3, pages 3–1513, 2001.
- [93] Mouser. LST1-01H06-GRN1-01. <https://mou.sr/45Lh6rl>, 2023.
- [94] Jurriaan D Mulder, Jack Jansen, and Arjen Van Rhijn. An affordable optical head tracking system for desktop vr/ar systems. In *Proceedings of the workshop on Virtual environments*, pages 215–223, 2003.
- [95] Kazuki Nakamura, Shingo Nakagawa, Hiroshi Matsubara, Daisuke Tatsui, Kiyotaka Seki, Shinichiro Haruyama, and Fumio Teraoka. Development of broadband telecommunications system for railways using laser technology. *Electrical Engineering in Japan*, 190(3):45–56, 2015.
- [96] S Narayana, R V Prasad, V Rao, L Mottola, and TV Prabhakar. Hummingbird: Energy efficient gps receiver for small satellites. In *Proc. of MobiCom*, pages 1–13, 2020.
- [97] T Nielsen, M Fricke, D Hellweg, and P Andresen. High efficiency beam splitter for multifocal multiphoton microscopy. *Journal of microscopy*, 201(3):368–376, 2001.
- [98] NIST. Single exponential smoothing. <https://www.itl.nist.gov/div89x8/handbook/pmc/section4/pmc431.html>, 2023.
- [99] NXP. MMA8452Q. <https://www.nxp.com/docs/en/data-sheet/MMA8452Q.pdf>, 2023.
- [100] Chin Wan Oh, Zizheng Cao, Eduward Tangdionga, and Ton Koonen. Free-space transmission with passive 2D beam steering for multi-gigabit-per-second per-beam indoor optical wireless networks. *Optics Express*, 24(17):19211–19227, 2016.
- [101] Hiromasa Oku and Masatoshi Ishikawa. High-speed liquid lens with 2 ms response and 80.3 nm root-mean-square wavefront error. *Applied Physics Letters*, 94(22):221108, 2009.
- [102] OpenMV. OV5640. <https://openmv.io/products/ov5640-fpc-camera-module>, 2023.
- [103] Edmund Optics. 22-827. <https://www.edmundoptics.com/p/132mm-fl-m12-mount-180-deg-fisheye-lens/49590/>, 2023.
- [104] Edmund Optics. 62-274. <https://www.edmundoptics.com/p/23quot-format-c-mount-fisheye-lens-18mm-fl/16922/>, 2023.
- [105] Edmund Optics. 67-119. <https://www.edmundoptics.com/p/532nm-25mm-diameter-od-4-notch-filter/21656/>, 2023.
- [106] Edmund Optics. 83-107. <https://www.edmundoptics.com/p/160mm-fl-no-ir-cut-filter-f4-micro-video-lens/26288/>, 2023.
- [107] Oralite. M82. <https://www.orafol.com/en/americas/products/oralite-m82-aids-to-navigation-tape>, 2023.
- [108] Dr. Rüdiger Paschotta. Polarization of laser emission, Oct 2018.
- [109] JA Picazo-Bueno, M Trusiak, and V Micó. Single-shot slightly off-axis digital holographic microscopy with add-on module based on beamsplitter cube. *Optics Express*, 27(4):5655–5669, 2019.
- [110] Ling Qin, Ben Niu, Bao-Shan Li, Xiao-Li Hu, and Yong-Xing Du. High precision indoor positioning algorithm of single led lamp based on a-bayes. *Optik*, 241:167190, 2021.
- [111] Bharath Ramesh, Hong Yang, Garrick Orchard, Ngoc Anh Le Thi, Shihao Zhang, and Cheng Xiang. DART: distribution aware retinal transform for event-based cameras. *IEEE transactions on pattern analysis and machine intelligence*, 42(11):2767–2780, 2019.
- [112] Hamada Rizk, Yuma Okochi, and Hirozumi Yamaguchi. Demonstrating hitonavi- μ : A novel wearable lidar for human activity recognition. In *Proc. of MobiCom*, 2022.
- [113] Miguel Fabián Romero Rondón, Lucile Sassatelli, Ramón Aparicio Pardo, and Frédéric Precioso. Track: a multi-modal deep architecture for head motion prediction in 360° videos. In *IEEE International Conference on Image Processing*, pages 2586–2590, 2020.

- [114] John Rzasa. *Pointing, acquisition, and tracking for directional wireless communications networks*. PhD thesis, The University of Maryland, 2012.
- [115] Wee-Leong Saw, Hazem H Refai, and James J Sluss Jr. Free space optical alignment system using gps. In *Free-Space Laser Communication Technologies XVII*, volume 5712, pages 101–109, 2005.
- [116] Qijia Shao, Amy Sniffen, Julien Blanchet, Megan E. Hillis, Xinyu Shi, Themistoklis K. Haris, Jason Liu, Jason Lamberton, Melissa Malzkuhn, Lorna C. Quandt, James Mahoney, David J. M. Kraemer, Xia Zhou, and Devin Balkcom. Teaching american sign language in mixed reality. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 4(4), dec 2020.
- [117] Qijia Shao, Jian Wang, Bing Zhou, Vu An Tran, Gurunandan Krishnan, and Shree Nayar. N-euro predictor: A neural network approach for smoothing and predicting motion trajectory. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 7(3), 2023.
- [118] K P Shaw and S Y Hsu. Horizontal distance and height determining falling pattern. *Journal of Forensic Sciences*, 43(4):765–771, July 1998.
- [119] Josef Sofka, Vladimir V Nikulin, Victor A Skormin, David H Hughes, and David J Legare. Laser communication between mobile platforms. *IEEE transactions on Aerospace and Electronic Systems*, 45(1):336–346, 2009.
- [120] Elahe Soltanaghaei, Avinash Kalyanaraman, and Kamin Whitehouse. Multipath Triangulation: Decimeter-Level WiFi Localization and Orientation with a Single Unaided Receiver. In *Proc. of MobiSys*, page 376–388, 2018.
- [121] In Keun Son and Shiwen Mao. A survey of free space optical networks. *Digital communications and networks*, 3(2):67–77, 2017.
- [122] Thad Starner, Bastian Leibe, David Minnen, Tracy Westyn, Amy Hurst, and Justin Weeks. The perceptive workbench: Computer-vision-based gesture tracking, object tracking, and 3D reconstruction for augmented desks. *Machine Vision and Applications*, 14:59–71, 2003.
- [123] Daria Stepanova, Valentin Pryanichnikov, Sergey Khandorin, Andrey Kulchitsky, and Alexey Kuznetsov. Developing a highly accurate pointing system for free space optical communications. *Journal of Communications*, 14(12), 2019.
- [124] Menlo Systems. APD210. <https://www.menlosystems.com/products/photodetectors/apd210/>, 2023.
- [125] Midwest Optical Systems. BN520. <https://midopt.com/filters/bn520/>, 2023.
- [126] Brahim Tamadazte, Rupert Renevier, Jean-Antoine Séon, Andrey V. Kudryavtsev, and Nicolas Andreff. Laser beam steering along three-dimensional paths. *IEEE/ASME Transactions on Mechatronics*, 23(3):1148–1158, 2018.
- [127] Thorlabs. Beamsplitter guide. https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=9028, 2023.
- [128] Thorlabs. C392TME-A. <https://www.thorlabs.com/thorproduct.cfm?partnumber=C392TME-A>, 2023.
- [129] Thorlabs. L638P. <https://www.thorlabs.com/thorproduct.cfm?partnumber=L638P200>, 2023.
- [130] Thorlabs. PBS25-633-HP. <https://www.thorlabs.com/thorproduct.cfm?partnumber=PBS25-633-HP>, 2023.
- [131] Thorlabs. S121C. <https://www.thorlabs.com/thorproduct.cfm?partnumber=S121C>, 2023.
- [132] Thorlabs. Small beam diameter scanning galvo mirror systems. https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=3770, 2023.
- [133] Thorlabs. TRS127-020-A. <https://www.thorlabs.com/thorproduct.cfm?partnumber=TRS127-020-A>, 2023.
- [134] Hideki Urabe, Shinichiro Haruyama, Tomohiro Shogenji, Shoichi Ishikawa, Masato Hiruta, Fumio Teraoka, Tetsuya Arita, Hiroshi Matsubara, and Shingo Nakagawa. High data rate ground-to-train free-space optical communication system. *Optical Engineering*, 51(3):031204, 2012.
- [135] Frank Van Diggelen and Per Enge. The world’s first GPS MOOC and worldwide laboratory using smartphones. In *Proceedings of the 28th International Technical Meeting of the Satellite division of the Institute of Navigation*, pages 361–369, 2015.
- [136] D Yu Velikovskiy, VE Pozhar, and MM Mazur. Acousto-optics devices for high-power laser beam. In *WDS*, volume 12, pages 65–68, 2012.
- [137] Vicon. Vicon documentation. <https://docs.vicon.com/pages/viewpage.action?pageId=107483982>, 2023.

- [138] Allied Vision. Alvium 1800 U240. <https://www.alliedvision.com/en/products/alvium-configurator/alvium-1800-u/240/>, 2023.
- [139] Anjitha Viswanath, Shailesh Singh, V. K. Jain, and Subrat Kar. Design and implementation of MOEMS based ground to satellite free space optical link under turbulence condition. *Procedia Computer Science*, 46:1216–1222, 12 2015.
- [140] Michele Volpi and Devis Tuia. Dense semantic labeling of subdecimeter resolution images with convolutional neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, 55(2):881–893, 2017.
- [141] Shane M. Walsh, Skevos F. E. Karpathakis, Ayden S. McCann, Benjamin P. Dix-Matthews, Alex M. Frost, David R. Gozzard, Charles T. Gravestock, and Sascha W. Schediwy. Demonstration of 100 Gbps coherent free-space optical communications at LEO tracking rates. *Scientific Reports*, 12(1):18345, 2022.
- [142] Dingkan Wang, Connor Watkins, and Huikai Xie. MEMS Mirrors for LiDAR: A Review. *Micromachines*, 11(5), 2020.
- [143] Greg Welch and Gary Bishop. An introduction to the kalman filter. *Proc. Siggraph Course*, 8, 01 2006.
- [144] Thomas Weyrauch and Mikhail Vorontsov. Free-space laser communications with adaptive optics: Atmospheric compensation experiments. *Journal of Optical and Fiber Communications Reports*, 1:355–379, 05 2010.
- [145] Linghui Yang, Yuanlin Pan, Jiarui Lin, Yang Liu, Yue Shang, Shuo Yang, and Hanwen Cao. Automatic guidance method for laser tracker based on rotary-laser scanning angle measurement. *Sensors*, 20(15):4168, 2020.
- [146] Liangchen Ye, Gaofei Zhang, Zhen You, and Chi Zhang. A 2D resonant MEMS scanner with an ultra-compact wedge-like multiplied angle amplification for miniature LIDAR application. In *Sensors, 2016 IEEE*, pages 1–3, 2016.
- [147] Lin Yongbing, Zhang Guoxiong, and Li Zhen. An improved cat’s-eye retroreflector used in a laser tracking interferometer system. *Measurement Science and Technology*, 14(6):N36, May 2003.
- [148] Yuan Zhuang, Jun Yang, You Li, Longning Qi, and Naser El-Sheimy. Smartphone-based indoor localization with bluetooth low energy beacons. *Sensors*, 16(5), 2016.
- [149] Jan Zizka, Alex Olwal, and Ramesh Raskar. Speckle-Sense: Fast, Precise, Low-Cost and Compact Motion Sensing Using Laser Speckle. In *Proc. of UIST*, page 489–498, 2011.
- [150] Mo Zohrabi, Robert H Cormack, and Juliet T Gopinath. Nonmechanical beam steering using tunable lenses. In *2017 Conference on Lasers and Electro-Optics (CLEO)*, pages 1–2, 2017.
- [151] Mo Zohrabi, Wei Yang Lim, Robert H Cormack, Omkar D Supekar, Victor M Bright, and Juliet T Gopinath. Lidar system with nonmechanical electrowetting-based wide-angle beam steering. *Optics Express*, 27(4):4404–4415, 2019.