

OWL: Enabling Orientation-Diverse Wireless Localization Research

Michael Sedlak, Jesús Pestana, Carlo Alberto Boano

Institute of Technical Informatics

Graz University of Technology

Graz, Austria

Email: {sedlak, pestana, cboano}@tugraz.at

Abstract—Device orientation is a known source of variability in wireless ranging and localization measurements, yet often underexplored in experimental evaluations. As we show in our empirical study using off-the-shelf Bluetooth Channel Sounding hardware, its impact can be substantial: systematically rotating a device over the full 360° range reveals differences up to $10\times$ in the median localization error. Unfortunately, most evaluations in the literature tend to fix device orientation or sample it only coarsely, and tools supporting fine-grained automated sweeps of device orientation remain scarce. To address this gap, we present the OWL, a cost-effective tool enabling automated wireless measurements with fine-grained control of device orientation with sub-degree accuracy. We demonstrate the depth of analysis unlocked by the OWL by collecting high-resolution datasets capturing the real-world performance of Bluetooth Channel Sounding as a function of device orientation. We further show that these high-resolution datasets allow us to attribute orientation-dependent ranging errors to two contributing factors: direct-path signal power—shaped by antenna-pattern gain and polarization mismatch—and algorithmic artifacts in the ranging implementation. We complete this work by openly releasing all our hardware, firmware, and datasets to lower the barrier to reproducible, orientation-diverse research.

Index Terms—Bluetooth 6.0; Channel Sounding; Location awareness; Ranging; Antenna radiation pattern; Polarization.

I. INTRODUCTION

As location-awareness gains importance in Internet-of-Things (IoT) deployments, modern applications increasingly require high-accuracy ranging and localization on low-cost, battery-powered, and resource-constrained devices [1]. *Localization* estimates the position of a mobile device (*tag*) from distance measurements to multiple fixed reference nodes with known positions (*anchors*) through multilateration. These distance measurements, referred to as *ranging*, rely on mechanisms such as time-of-flight (ToF) or phase-based ranging (PBR), and their accuracy depends on the underlying technology. For instance, Ultra-wideband (UWB) radios allow for accurate ToF estimation, yielding a mean absolute error (MAE) below 6 cm in line-of-sight (LoS) conditions [2]. Bluetooth Low Energy (BLE) channel sounding (CS), introduced in the Bluetooth Core Specification 6.0 [3], combines round-trip time and PBR [4], [5], with current baseline implementations achieving a root mean square error (RMSE) of 1.04 m [6]. Crucially, these reported accuracies are typically achieved under favorable conditions, particularly w.r.t. device orientation. Indeed, tag orientation is rarely controlled and can drastically affect ranging and localization accuracy [7]–[9].




Fig. 1. Location estimates color-coded as a function of the tag’s orientation. The tag performs a 360° rotation in 1° steps; 0° corresponds to the tag facing left (along the negative x -axis). Two clusters and their corresponding CDFs are highlighted: a favorable tag orientation window (-126° to -106° , median error of 63 cm) and an unfavorable one (79° to 99° , median error of 6.30 m).

Fig. 1 illustrates this for an indoor BLE CS localization system (whose setup is described in § IV-A) operating in LoS conditions: a fine-grained 360° tag orientation sweep produces location estimates (color-coded by tag orientation) that form distinct spatial clusters, exposing a pronounced dependence of localization accuracy on device orientation. For instance, a favorable orientation window (-126° to -106° in Fig. 1) yields a median localization error of 63 cm, while an unfavorable window increases it to 6.30 m, corresponding to a $10\times$ increase (see CDFs in Fig. 1). Furthermore, the 99th percentile (p99) localization error increases by an even larger factor of $22.5\times$ (from 0.88 m to 19.84 m). These results underscore that device orientation is a critical factor; evaluations that neglect it risk to *substantially misrepresent real-world performance*.

Lack of orientation diversity in current evaluations. Despite this evidence, most existing datasets and experimental evaluations of wireless localization performance fix device orientation, typically positioning tags and anchors face-to-face with intentionally aligned antennas [10]–[13]—so to establish a *best-case* baseline and eliminate orientation as a confounding factor—and focus on isolating specific challenges such as the presence of non-line-of-sight (NLoS) conditions [10].

Lack of suitable measurement tools. This limited consideration of device orientation diversity is largely driven by the practical challenges and significant overhead associated with precisely, systematically, and reproducibly controlling device orientation during measurement campaigns. Because of this, experimental studies are often restricted to a small set of

discrete orientations (e.g., 45° intervals) [7], [8], [14], [15], which risk missing effects that manifest over narrow angular ranges [9]. Conversely, studies relying on automated device orientation sweeps employ motion-tracked robots [16]–[22], spherical near-field scanners [23], or custom motorized platforms [24]–[26]: all these setups are, however, either expensive, proprietary, or insufficiently documented for replication. As a result, the sometimes gradual, sometimes abrupt signal transitions arising from the interplay between the environment, multipath propagation, device orientation, as well as antenna polarization and radiation patterns remain undersampled across the literature. Beyond ranging, the lack of accessible, precise rotational tools also hinders research exploring the Angle-of-Arrival (AoA) capabilities of modern radios [16], [24]; more generally, it hinders the study of orientation-induced effects on wireless communication and sensing.

Contributions. To tackle these limitations and democratize experimentation with fine-grained orientation diversity, we present the OWL (Orientation-aware Wireless Localization), an open-source, cost-effective tool designed to automate the acquisition of wireless measurements with sub-degree angular accuracy, and demonstrate it in action, highlighting the depth of analysis it enables. Our core contributions are:

- **Design of a practical measurement tool (§ II).** We build a measurement platform that leverages an absolute angle encoder to enable an automated and fine-grained control of device orientation. We make our hardware design, fabrication files, and control software openly available.
- **Tool accuracy and repeatability evaluation (§ III).** We characterize the OWL’s angular accuracy and repeatability, obtaining a mean RMSE of 0.38° across two builds.
- **Orientation-aware BLE CS dataset (§ IV).** We use the OWL to collect a densely-sampled (1° -steps) dataset capturing BLE CS performance as a function of device orientation across four real-world environments. Our data reveals variations in the ranging and localization error by up to tens of meters across different device orientations.
- **In-depth analysis of orientation impact (§ V).** We leverage the datasets’ high sampling density to attribute the degraded ranging performance to reduced direct-path signal power and to uncover BLE CS algorithmic artifacts.
- **Discussion of future research opportunities (§ VI).** We finally discuss how the OWL can support future research in wireless communication, localization, and sensing.

II. INSIDE THE OWL

The OWL is an open-source, cost-effective, single-axis rotation platform designed to automate the collection of wireless measurements over fine-grained device orientations.

Requirements. The design of the OWL is guided by four key requirements. (i) *Dependable rotation*: the platform must guarantee automated, sub-degree angular positioning accuracy across the full 360° range without mechanical drift or cable tangling, ensuring repeatability across multiple runs and independent builds. (ii) *Affordability and accessibility*: the bill of




Fig. 2. Key components of the OWL tool with a mounted radio device.

materials must rely on inexpensive, off-the-shelf components, and standard rapid-prototyping techniques to enable low-cost replication. (iii) *Ease of use*: the platform must require minimal technical expertise and effort to set up and operate. (iv) *Flexibility and openness*: hardware (HW) and firmware must be both open-source and payload-agnostic, in order to allow other researchers to easily use and extend the platform.

Design. Fig. 2 shows the OWL design: all cabling is routed through a hollow-shaft brushless direct current (BLDC) motor via a slip ring. This design enables endless 360° rotation without mechanical interference from the cabling. A built-in USB hub allows the host computer to communicate with both the OWL and the payload through a single USB connection. We use an MPS MA600A magnetic absolute angle encoder [27] to provide angular position feedback. This contactless sensor measures the orientation of a diametrically magnetized permanent magnet attached to the motor shaft, it is immune to mechanical wear, and exhibits negligible drift—yielding sub-degree accuracy and repeatable performance while providing an absolute angular position reading at power-on. The motor is controlled via the open SimpleFOC library [28], configured with a maximum speed of $180^\circ/\text{s}$ and a cascade position controller: a proportional (P) loop for position and a proportional-integral (PI) loop for velocity regulation.

To address the ease-of-use requirement, the control board features a plug-and-play USB interface that presents a standard virtual COM port, enabling communication via a simple, human-readable serial text protocol without requiring custom drivers or specialized host-side software. The system is powered via the USB Power Delivery standard and includes onboard status LEDs to provide direct visual feedback. Furthermore, the design incorporates SparkFun Qwiic connectors [29] for solderless I²C expansion, facilitating the integration of auxiliary sensors. Together with the publicly released designs, the payload-agnostic USB pass-through and a generic mount allow researchers to attach arbitrary radios or sensors, satisfying the flexibility requirement.

To ensure affordability, the OWL uses consumer-grade, off-the-shelf HW available from established vendors, custom printed circuit boards (PCBs) from low-cost prototyping services, and 3D-printed structural components, resulting in a cost of approximately € 170 per unit (including shipping and taxes, at an order quantity of five units) as of mid-2026 — one to three orders of magnitude below hardware used in prior work.

Synchronization and data logging. During an orientation sweep, the host computer instructs the OWL to rotate to a target angle via serial communication. The OWL monitors both orientation error and rotational speed; once both fall below configurable thresholds, it notifies the host that the target angle has been reached. The host then directly triggers the radio measurement on the payload through the USB pass-through connection. Data logging, including the angle, timestamp, and radio measurements, is managed by the host computer.

Open-source artifacts¹. To foster reproducibility and community adoption, we openly release the OWL’s HW design, 3D-printable models, control firmware, and collected datasets.

III. OWL VALIDATION

We validate the angular positioning performance of two OWL builds against a sub-millimeter-accurate Qualisys motion-capture (mocap) system serving as a ground-truth reference. We characterize two complementary aspects of this performance: (i) the *angular accuracy*, i.e., how closely the platform reaches a given target angle (relative to its own zero) across the full 360° range, and (ii) the *per-position repeatability*, i.e., the precision with which it returns to each individual target angle under a shuffled sequence of non-consecutive positions.

Ground-truth reference. The mocap system serves as our tracking baseline. A cross-shaped constellation of passive reflective markers is rigidly attached to the top of the OWL’s rotating platform, with markers placed at the endpoints and center of four roughly perpendicular arms forming an elliptical shape with a major and minor axis (approximately 136 mm and 94 mm across). This asymmetric arrangement forms a rigid body whose orientation is tracked in real time. Before each experiment, the mocap system is calibrated so that the initial orientation of the OWL defines the zero reference. Conveniently, the system’s baseline tracking uncertainty can be derived directly from the calibration results provided by the mocap system. After wand calibration, the mocap system reports a standard deviation of the wand length of approximately $\sigma_L \approx 0.71$ mm, which translates to a per-marker position uncertainty of $\sigma_{\text{marker}} = \sigma_L / \sqrt{2} \approx 0.5$ mm (via propagation of uncertainty [30]). Given the radial distances r_i of each marker from the OWL’s center of rotation, the resulting angular tracking uncertainty is $\sigma_\theta = \sigma_{\text{marker}} / \sqrt{\sum r_i^2} \approx 0.25^\circ$, well below the sub-degree accuracy target of the OWL.

Setup. To evaluate the angular positioning accuracy and repeatability of the OWL, we perform two distinct experiments on each of the two builds (1 and 2). Both experiments automate measurements following the synchronization and data logging mechanism described in § II. First, in a consecutive rotation sweep (E_1), the platform is rotated in 1° incremental steps over a full 360° rotation. We record 10 ground truth measurements from the mocap system per step, and repeat this for 2 runs per build. Second, in a shuffled target-angle sweep (E_2), the platform is instructed to visit all 360 target angles in a randomly shuffled order across 5 distinct runs per build. We




Fig. 3. Accuracy and repeatability of two OWL builds: (a) empirical CDFs, RMSE, and p99; (b) signed angular error versus the target angle (solid lines show the mean, shaded regions represent the min-max range over five runs).

use the same five randomized sequences across both builds to ensure a fair comparison. At each target angle, the angular error is computed as the difference between the OWL-reported angle and the mocap-measured ground-truth.

Angular accuracy. Fig. 3(a) shows the combined cumulative distribution function (CDF) of the absolute angular errors across all samples from both experiments. Both builds achieve sub-degree angular accuracy, with RMSE values of 0.38° and 0.39°. Furthermore, the error is tightly bounded: the 99th percentile (p99) of the error remains below 1° for both builds.

Per-position angular repeatability. To assess how consistently the OWL reaches each target angle when rotating between non-consecutive positions, we analyze the signed angular error during the shuffled sweep (E_2) in detail. Fig. 3(b) plots the mean error (solid lines) and the min-max error envelope (shaded regions) across the five independent runs for all 360 target angles. The figure legend reports two complementary error-spread metrics. The error remains mostly within $\pm 1^\circ$ across the full sweep, with the maximum per-target-angle spread (i.e., the worst-case variation at any single angle across the five runs) ranging from 0.70° to 0.72°, while the overall error spread across all angles and runs remains within 1.25° to 1.54°. These small error spreads are sufficient to resolve the orientation-induced effects we observe, whose transitions unfold over much wider angular scales (§ IV + § V).

Key takeaways. The evaluation across two independently assembled OWL builds shows that the platform achieves sub-degree angular positioning accuracy and repeatability. Importantly, these results establish a reproducible performance baseline: researchers who build their own OWL can reasonably expect it to perform within the bounds characterized here, without the need for an external motion-capture system for calibration or verification. The OWL thus provides a cost-effective, accessible tool for automated, high-accuracy orientation sweeps in wireless communication and sensing research.

IV. OWL IN ACTION

We now leverage the OWL to investigate orientation-induced effects in BLE CS localization and ranging. To this end, we conduct automated 360° device orientation sweeps in line-of-sight (LoS) conditions across several real-world environments using commercial off-the-shelf hardware. We first describe the experimental setup used to perform our measurement campaign (§ IV-A), and then analyze the collected datasets,

¹Open-source artifacts available at: <https://iti.tugraz.at/OWL>.

highlighting the high variability introduced by device orientation (§ IV-B). The underlying causes of this orientation-dependent variability are analyzed in detail in § V.

A. Experimental Methodology and Setup

Hardware platform. We use the Nordic Semiconductor nRF54L15-DK, which fully supports BLE CS and features an inverted-L PCB antenna approximately one-quarter wavelength long. The device runs the nRF Connect SDK (v3.2.4) on the Zephyr operating system and transmits at 0 dBm [31]. We enable the full CS channel map (72 channels not overlapping with BLE’s advertising channels) and employ a channel map repetition of one. Unless otherwise stated, all remaining parameters are kept at vendor-default values. Our firmware is based on the Nordic SDK channel sounding examples², with modifications to log IQ samples, tone quality indicators, and per-tone received signal strength indicator (RSSI)—derived from the IQ amplitude and reference power level of each ranging sub-event—in a machine-readable JSON format. We additionally remove the default median filter over the last nine ranging samples to obtain raw, unfiltered estimates.

Real-world environments. Experiments are conducted under LoS conditions in three outdoor (PARK, YARD, and STREET) and one indoor (LECTURE HALL) environments. LECTURE HALL is an 88 m² university room with a ceiling height of 3.25 m; it contains rows of tables (73 cm height) and chairs. PARK and YARD feature relatively flat grass terrain, while STREET has a paved surface with a slight ramp. The outdoor areas are generally free of large reflective structures; the nearest significant reflectors are a row of greenhouses approximately 6 m from the measurement axis in YARD, and a building wall and metal fence approximately 9 m away in STREET. In PARK, the closest building is approximately 20 m away, with some trees and construction materials at around 8 m. Initiator and reflector are mounted on tripods at a height of 1.4 m outdoors and 1.9 m in LECTURE HALL (to ensure LoS clearance above tables). Ground-truth distances outdoors are measured with a surveying tape (cm-level accuracy); in LECTURE HALL, positions are determined using a total station. To reduce systematic ranging biases, we calibrate each initiator-reflector pair per environment and setup by subtracting a per-pair offset, defined as the median ranging error across all orientations and positions.

Orientation sweep. We use the OWL to rotate the BLE CS initiator clockwise through a full 360° range in 1° steps, and henceforth refer to the resulting angular position as the *OWL angle*. Fig. 4 shows the rotation direction and the 0° reference used in our setup. We employ two different settings: a *vertical* setup (Fig. 4a)—representing a typical deployment orientation in which both antenna-pattern gain and polarization effects jointly influence the received signal—and a *horizontal* setup (Fig. 4b)—in which the antenna-pattern gain is expected to remain approximately constant throughout the OWL’s rotation.




Fig. 4. Device orientations explored in our measurement campaigns. The BLE CS initiator is mounted on top of the OWL, which rotates clockwise in the direction of the arrow. The reflector device remains static.

Outdoors, we perform two independent measurement runs at fixed distances of 10, 20, and 25 m between one initiator and one reflector. Indoors, we perform three independent runs with four reflectors serving as anchors (marked as red triangles in Fig. 1) positioned in a rectangle around the initiator (tag, marked as a yellow star). Distances between tag and anchors 1–4 are 4.22, 4.46, 4.28, and 4.14 m, respectively.

B. Impact of Orientation on Localization and Ranging

Device orientation affects both ranging and localization. Fig. 1 shows the localization error in LECTURE HALL, which exhibits a clear dependence on the OWL angle, with clusters following the angular color coding. Accurate localization is achieved only for a narrow range of device orientations, with the median localization error varying from 63 cm in favorable orientations to 6.3 m in unfavorable ones, a tenfold difference (§ I).

This high variability in the median localization error stems from the underlying variability of ranging estimates as a function of device orientation. Fig. 5 shows the absolute ranging error as a function of the OWL angle across all outdoor and indoor vertical setup experiments. Examining the per-anchor ranging errors in LECTURE HALL, the lowest errors are generally found in the -90° to 0° range for all four anchors. Parts of the angular sweep show a gradual, low-spread variation of the ranging error, revealing a clear deterministic dependence on the rotation angle. In regions where the error jumps abruptly, the spread of the estimates is also large. Notably, Anchor 1 achieves considerably lower ranging error than the other three anchors, excluding one large spike, despite all anchors being at similar distances (≈ 4 – 4.5 m) under LoS conditions. We attribute these discrepancies to multipath propagation effects, which are common in indoor environments.

To isolate the impact of device orientation from multipath effects, we turn to open outdoor environments. In PARK, YARD, and STREET, the error is consistently low around 0° (roughly from -90° to 45° – 75°) and around $\pm 180^\circ$ (from -135° to 135° passing through $\pm 180^\circ$). Conversely, two angular bands (approximately from -135° to -90° and from 75° to 135°) exhibit systematically higher errors across all three environments and investigated distances. These bands appear at consistent angles across environments and runs, indicating a deterministic, orientation-dependent effect. To quantify this, we slide a 45° orientation window ($\pm 22.5^\circ$) per configuration and compare the p99 ranging error at the best- and worst-case orientations. For the best-case orientations, the p99 error stays below 0.58 m; for the worst-case ones, it reaches up to 18.8 m (at 25 m in PARK), i.e., we observe up to $37\times$ differences caused by changes in device orientation.

²Samples: channel_sounding_ras_initiator and_reflector.




Fig. 5. Absolute BLE CS ranging error vs. OWL angle for all outdoor and indoor experiments.

V. IN-DEPTH ANALYSIS OF ORIENTATION IMPACT

The results of § IV establish that the BLE CS ranging performance varies with device orientation, but not *why*. Leveraging the OWL’s fine-grained sweeps, we investigate two contributing factors: the direct-path signal power (§ V-A) and algorithmic artifacts in the PBR ranging estimation (§ V-B).

A. Direct-Path Signal Power

The orientation of a device’s antenna relative to its communication peer affects the received signal power through two mechanisms: the directional gain of the *antenna radiation pattern* and the *polarization alignment* between transmitter and receiver. The vertical setup of § IV-A exposes both effects jointly. To remove the influence of the antenna characteristic and isolate the polarization component, we additionally introduce a *horizontal setup* (Fig. 4b). Here, the antenna-pattern gain toward the peer device remains approximately constant during rotation, so RSSI variations are predominantly due to transitions between co- and cross-polarization. We collect horizontal setup sweeps in the three outdoor environments following the same procedure used for the vertical setup (§ IV-A). Fig. 6 plots the median per-tone RSSI as a function of the OWL angle for both the vertical and horizontal setups.

Vertical setup (combined effects). The measured data reveal two prominent RSSI dips, separated by roughly 200° – 240° rather than a symmetric 180° . We attribute this to two overlapping effects: directional antenna gain and polarization mismatch. These angular RSSI variations line up with the outdoor ranging errors shown in Fig. 5. The $[-90^\circ, 0^\circ]$ interval consistently exhibits the highest RSSI within each sweep and, correspondingly, the lowest ranging error across all outdoor setups, whereas the regions around $\pm 90^\circ$ show generally low RSSI and correspondingly high ranging error. This highlights a clear link between low RSSI and high ranging error.

Horizontal setup (polarization dominant). Our data exhibit two RSSI dips separated by $\approx 180^\circ$, matching the theoretical expectation for cross-polarization nulls. This is consistent with polarization mismatch being the dominant effect in this configuration. However, the dip positions are shifted by $\approx 40^\circ$ from the expected $\pm 90^\circ$ locations. We attribute this offset primarily to ground reflections, uneven terrain, and residual manual misalignment. The angular tone RSSI variation in the horizontal setup ranges from approximately 10.3 dB to 17.1 dB across configurations, smaller than the variation observed in the vertical setup (10.7 dB to 23.2 dB).




Fig. 6. Median tone RSSI vs. target angular position for vertical and horizontal setups in all three outdoor environments.

B. Algorithmic Artifacts in PBR Ranging

Having linked low RSSI to high ranging error, we now examine where these errors originate, tracing them to how the ranging algorithm processes the measured channel into a distance estimate. Distance estimation is performed by *Nordic’s open-source SDK PBR implementation* [31]. PBR estimates distance by measuring the wireless channel between an *initiator* and a *reflector* through the exchange of unmodulated tones across different frequencies in the 2.4 GHz band. Both devices capture the in-phase (I) and quadrature (Q) components of the received signals. To compute the range, these per-channel IQ samples from initiator and reflector are combined via complex multiplication, yielding a sampled estimate of the two-way channel frequency response (CFR). The CFR is then zero-padded and converted to the time domain using an Inverse Fast Fourier Transform (IFFT), producing a channel impulse response (CIR) whose peak locations correspond to propagation delays and, thus, distances. The CIR is sampled at discrete points or *bins*. The finite bandwidth of the BLE CS channel map limits the CIR resolution in two ways. First, each peak has a minimum mainlobe width, so closely spaced multipath components merge into a single broadened peak rather than appearing as separable peaks. Second, the implicit rectangular windowing of the IFFT surrounds each peak with sidelobes, i.e., smaller spurious peaks at fixed offsets from the true one.

Identifying CIR classes. To explain how CIR distortions affect ranging accuracy, we use representative samples from a single sweep in STREET (vertical setup) to illustrate *CIR classes* (Fig. 7). Each sample is marked at its angular position




Fig. 7. Visualization of algorithmic artifacts in Nordic's IFFT-based PBR implementation. Top row: full-scale and zoomed ranging-error plots, and median tone RSSI vs. OWL angle. Middle and bottom rows: relative CIR amplitude vs. corresponding distance for representative ranging samples mapping to specific points in the top plots. On each CIR, the marker is placed at the estimated distance; the dashed red line marks the true distance (25.0 m).

in the top-row plots and linked to its individual CIR and distance estimate in the bottom plots. We begin with the accurate classes in the high-RSSI regions (circles, e.g., the orange one in Fig. 7c), which feature relatively undistorted CIRs, then focus on the attenuated regions where the high-error classes, each arising from a specific type of CIR distortion, are more prominent (triangles and turquoise squares).

Ranging accuracy mainly depends on whether the direct-path peak is visible and correctly selected. To find the direct-path peak, the algorithm starts from the strongest peak in the CIR and searches leftward, toward shorter distances, to favor the first-arriving path over later multipath. It selects the earliest peak whose amplitude is above 40% of the maximum, then measures that peak's width down to its left null (i.e., where it falls below 10% of its own amplitude). The algorithm compares this width against a threshold set by the minimum mainlobe width (>7 bins), selecting one of two branches.

If the selected peak is narrow, it is consistent with a single resolved path and its position is refined via closed-form parabolic interpolation to sub-bin resolution. Four CIR classes arise: (1) **Strong direct path (Strong DP)**: the direct-path peak is itself the dominant peak, yielding high accuracy (orange circle; Fig. 7d,e). (2) **Weak direct path (Weak DP)**: the direct-path peak is not the dominant peak but still exceeds the 40% threshold, so parabolic refinement still yields a low-error estimate. This covers both a clearly weak direct path (green circle; Fig. 7e,f) and a strong one that a later multipath peak exceeds only slightly (blue circle; Fig. 7d). (3) **Suppressed direct path (Supp DP)**: the direct-path peak is absent or below the threshold, so the algorithm locks onto a later multipath

peak and *overestimates* the distance (red triangle; ranging error visible in Fig. 7a, CIR in Fig. 7e). (4) **IFFT sidelobe (Sidelobe)**: a windowing sidelobe appears at a shorter distance and is mistaken for the direct-path peak, causing the algorithm to *underestimate* the distance (purple triangle; Fig. 7f). Because the sidelobe sits at a fixed offset from a propagation-path peak, this underestimation is consistent across a range of angles and appears as a band spanning 130° to -130° , i.e., around $\pm 180^\circ$, in the orientation sweep (Fig. 7a).

If the selected peak is wide, multiple paths have overlapped and broadened the direct-path peak into what we term a *broad direct path* (BDP). Its maximum occurs later than the true arrival, so using the peak position would overestimate the distance. Instead, the algorithm takes the BDP's left null plus a fixed offset (half the minimum mainlobe width, 7 bins). Because both the left null and the offset are measured in whole IFFT bins, the distance estimates can only take discrete values spaced ≈ 29.3 cm apart, producing the characteristic horizontal bands in the orientation sweep (squares; Fig. 7b). This yields two CIR classes: (1) **Low-error BDP (LoErr BDP)**: the left-null branch returns an estimate close to the true distance (brown square; Fig. 7g). (2) **High-error BDP (HiErr BDP)**: in the example (turquoise square; Fig. 7g), a sidelobe overlaps with the downward slope of the BDP, so its left null is detected too early, yielding a large ranging error.

Impact and prevalence of CIR classes. Using the six class labels defined above, we label every ranging estimate in the outdoor datasets. We assign each sample's class by combining the algorithm's internal branching decisions (BDP vs. non-BDP), the ground-truth distance, and whether the estimate is an inlier. To separate inliers from outliers, we compute the median absolute deviation (MAD) of the ranging errors following [7]. Under the assumption of a Gaussian inlier distribution with a long outlier tail, the MAD yields an estimated inlier standard deviation of $\hat{\sigma} = 1.4826 \cdot \text{MAD}$. We set the inlier/outlier threshold at $3\hat{\sigma}$ (1.16 m). This threshold serves as a surrogate to verify whether the direct-path peak was selected, and to distinguish low-error (Strong DP, Weak DP, and LoErr BDP) from high-error classes (Supp DP, Sidelobe, and HiErr BDP). Finally, the sign of the ranging error is used to disambiguate the remaining cases: distinguishing overestimation (Supp DP) from underestimation (Sidelobe).

The accurate classes (Strong DP, Weak DP, and LoErr BDP) account for 86.58% of samples and have high mean tone RSSI (-69.6 dBm for Strong DP). By contrast, the high-error classes (Supp DP, Sidelobe, and HiErr BDP) comprise the remaining 13.42% and suffer from severe degradation: their median ranging errors jump to 2.2–2.8 m ($>4\times$ larger than the 0.15–0.54 m of accurate classes) and their p99 errors jump to 12.6–14.9 m ($>10\times$ larger). High-error classes predominantly occur in low-RSSI regimes (mean tone RSSI from -80.2 to -76.7 dBm) and are *highly orientation-dependent* (Fig. 8): their relative frequency reaches $\geq 30.0\%$ in the ranges $[-130^\circ, -112^\circ]$ and $[49.0^\circ, 55^\circ]$, but drops below 5.0% in the ranges $[-87^\circ, 6^\circ]$ and $[148^\circ, 160^\circ]$ (i.e., those with high RSSI, see Fig. 5). This




Fig. 8. Relative frequency distribution of CIR classes vs. OWL angle. The dotted red line (HiErr Freq) marks the boundary separating low-error labels (bottom three categories) from high-error labels (top three categories).

$6\times$ variation in their prevalence confirms that orientation-induced signal attenuation and polarization mismatch drive the transition to high-error CIR classes. While our CIR classes are tied to Nordic’s IFFT-based implementation, we expect the underlying causes to generalize: direct-path suppression and peak broadening affect distance estimation in band-limited channels, while orientation-driven signal-power variations predominantly affect single-antenna BLE CS hardware. Taken together, our results demonstrate that the OWL’s automated sweeps can expose orientation-induced effects (from antenna-pattern and polarization signatures to algorithm branching behavior) that would be difficult to identify, let alone characterize, with manual or coarse-grained rotation methods.

VI. DISCUSSION AND FUTURE WORK

Our results confirm that device orientation largely degrades BLE CS ranging performance, revealing orientation-dependent RSSI variations of up to 23.2 dB that are compounded by ranging algorithm artifacts. When propagated through multilateration, these ranging errors cause a tenfold increase in median indoor localization error, as shown in Fig. 1. These findings motivate several directions for future work.

Cross-technology benchmarking. Whilst our study focuses on BLE CS, the OWL and the methodology used in this work are generic and technology-agnostic, and can therefore be reused to benchmark and study the impact of device orientation across other wireless technologies (e.g., Wi-Fi, mmWave, and UWB) and transceiver generations (e.g., radios compliant to IEEE 802.15.4a vs IEEE 802.15.4z [8]). As an illustrative example, Fig. 9 presents UWB localization results obtained under the same indoor conditions used for BLE CS in Fig. 1. We use Qorvo’s DWM3001CDK devkit, configured with recommended settings, and perform double-sided two-way rangings on channel 9. Our results show that device orientation also impacts UWB performance: between the most favorable and unfavorable device orientation windows, the median and p99 localization errors increase by $2.2\times$ and $7.4\times$.

Adding NLoS conditions to the loop. While our experiments reveal substantial performance variability, they are conducted under LoS conditions, which are comparatively favorable. We expect non-line-of-sight (NLoS) conditions to further exacerbate orientation effects, introducing stronger multipath propagation and additional error sources. We will leverage the OWL to investigate these conditions in future work.

Enabling multi-axis rotation. Antenna patterns span azimuth and elevation, which single-axis rotation cannot capture. Future work will extend the OWL to enable multi-axis sweeps.




Fig. 9. UWB localization performance using the same setting as in Fig. 1. Two tag orientation windows and their respective CDFs are highlighted: a favorable one (110° , 130°) and an unfavorable one ($[-117^\circ, -97^\circ]$).

VII. RELATED WORK

Device orientation directly affects the characteristics of the wireless channel, e.g., by altering antenna gain, polarization, and phase response—a challenge well-documented by radio manufacturers. For instance, SiLabs [9] highlighted that polarization mismatches caused by changes in device orientation result in significant ranging bias (up to several meters) in single-antenna BLE CS setups, and recommended antenna-path diversity to mitigate these effects. Academic works confirm such orientation-induced impact, quantifying severe performance variability, e.g., in RSSI [15], [26].

Device orientation impact on ranging/localization. These orientation-induced physical-layer variations directly affect the ranging and localization performance of wireless technologies. For example, in UWB systems, ranging errors have been experimentally shown to vary from a 25 cm bias (caused by switching between ideal horizontal and vertical setups [8]) to ≈ 50 cm for continuous device rotations introduced using drones [32]. Similarly, for BLE CS, empirical studies have observed orientation-dependent ranging errors varying from mixed horizontal/vertical setup biases [14] to calibration offset shifts of up to 138 cm for single-antenna setups [7]. Device orientation was also shown to impact Wi-Fi localization [33]: ignoring it degrades Wi-Fi fingerprinting accuracy by 10–30% and increases localization error. To address these challenges, researchers rely on calibration models (e.g., inferring pose-dependent variance from first-path power [17]) or antenna-path diversity [7], [9]. Our proposed tool, the OWL, directly facilitates the development of such solutions by enabling an automated collection of large orientation-diverse datasets that can be used to train robust ranging and localization models. Even more, the fine-grained datasets collected by the OWL enable detailed analyses that can reveal weaknesses in algorithmic implementations, as we have shown in § V-B.

Inferring device orientation. Modern low-power radios support direction finding. For instance, Bluetooth 5.1 standardized the use of a constant tone to facilitate angle-of-arrival (AoA) and angle-of-departure (AoD) estimation. Similarly, new-generation UWB radios support Phase-Difference-of-Arrival (PDoA)-based measurements [16], [22]. Whilst many studies demonstrated the feasibility of AoA/AoD on off-the-shelf HW, they also highlighted how the accuracy of angular estimations

varies with device orientation. For example, Dotlic *et al.* [34] and Van Herbruggen *et al.* [22] have demonstrated high reliability only within a $\pm 80^\circ$ and $\pm 45^\circ$ angular window. ML models can be leveraged to maintain high AoA accuracy despite polarization mismatches and device orientation ambiguities [23], [24], but require high-resolution datasets. Our study is orthogonal to this body of research: the same methodology used to quantify/analyze orientation-induced variability in ranging and localization can likewise be applied to assess and address orientation-induced effects in AoA/AoD performance.

VIII. CONCLUSIONS

Device orientation is a major but often overlooked source of variability in wireless ranging and localization, causing up to 10 $\times$ differences in localization error. We have introduced the OWL, a cost-effective system enabling automated, fine-grained control of device orientation for reproducible measurement studies. Using the OWL, we have quantified BLE CS ranging and localization performance as a function of device orientation, attributed the high performance variability to antenna gain and polarization mismatch effects, and uncovered BLE CS algorithmic artifacts. By releasing our hardware, firmware, and datasets, we aim to enable and promote more rigorous, orientation-diverse wireless localization research.

ACKNOWLEDGMENTS

Funding agencies. This work was conducted within the SPIDR² project (“Secure, Performant, Intelligent, Dependable, Reliable, and Resilient Wireless Systems”) financed by the Technology Innovation Institute.

AI disclosure statement. Substantial portions of the source code, including OWL firmware, automation scripts, analysis utilities, and plotting scripts, were developed with assistance from Google Antigravity. The authors reviewed, integrated, tested, and modified the generated code and remain responsible for the artifacts and their conclusions. The authors have also used Google Gemini, ChatGPT, and Claude for linguistic refinement and grammar correction. All intellectual content and final revisions were managed entirely by the authors.

CRedit author statement. *M. Sedlak*: conceptualization, methodology, hardware design, software, investigation, formal analysis, data curation, visualization, validation, writing (original draft); *J. Pestana*: conceptualization, methodology, investigation, formal analysis, visualization, validation, writing (original draft); *C.A. Boano*: conceptualization, supervision, methodology, writing (reviewing & editing).

REFERENCES

- [1] F. Pereira *et al.*, “Challenges in Resource-Constrained IoT Devices: Energy and Communication as Critical Success Factors for Future IoT Deployment,” *Sensors*, vol. 20, no. 22, 2020.
- [2] E. Ember *et al.*, “Improving the Correction of NLoS-Induced Ranging Errors in UWB Systems through Enhanced Labeling,” in *Proc. of the 15th IFIP Wireless and Mobile Networking Conference (WMNC)*, 2024.
- [3] Bluetooth SIG, “Bluetooth Core Specification, Version 6.0,” Bluetooth® Technology Website, Aug. 2024, accessed: May 12, 2026.
- [4] Y. Schröder *et al.*, “Accurate and Precise Distance Estimation from Phase-Based Ranging Data,” in *Proc. of the 9th IPIN Conf.*, 2018.
- [5] Y. Schröder *et al.*, “Investigation of Multipath Effects on Phase-based Ranging,” in *Proc. of the 10th IPIN Conf.*, 2019.
- [6] A. Tsemko *et al.*, “Data-driven Processing using Parametric Neural Network for Improved Bluetooth Channel Sounding Distance Estimation,” in *Proc. of the IEEE 50th ICASSP Conf.*, 2025.
- [7] D. Dzafic *et al.*, “Orientation Matters: An Empirical Study on Ranging Performance of Bluetooth 6.0 Channel Sounding,” in *Proc. of the 23rd Int. Conf. Embedded Wireless Systems and Networks (EWSN)*, 2026.
- [8] T. Polonelli *et al.*, “Performance Comparison between Decawave DW1000 and DW3000 in low-power double side ranging applications,” in *IEEE Sensors Applications Symp. (SAS)*, 2022.
- [9] Silicon Laboratories Inc., “AN1493: Antenna Design Guidelines for BLE Channel Sounding,” Tech. Rep., 2025, rev. 0.3. Accessed: 2026-05.
- [10] M. Stocker *et al.*, “Performance of Support Vector Regression in Correcting UWB Ranging Measurements under LOS/NLOS Conditions,” in *Proc. of the CPS-IoT Worksh.*, 2021, dataset: <http://iti.tugraz.at/uwb-nlos>.
- [11] K. Bregar *et al.*, “Improving Indoor Localization Using Convolutional Neural Networks on Computationally Restricted Devices,” *IEEE Access*, vol. 6, 2018, eWINE localization dataset: <https://github.com/ewine-project/UWB-localization/tree/master/data/localization>.
- [12] M. Wang *et al.*, “A 600-m UWB Ranging Dataset With Radio Frequency Features and Deep Learning Applications,” *IEEE Sensors Journal*, vol. 25, 2025, <https://github.com/futurexlab2012/LR-UWB-Dataset>.
- [13] L. Flueratoru *et al.*, “Challenges in platform-independent UWB ranging and localization systems,” in *Proc. of the 16th ACM Worksh. on WINTeCH*, 2022, dataset: <https://zenodo.org/records/6984698>.
- [14] J. Wieme *et al.*, “Performance Evaluation of Bluetooth Channel Sounding on Commercial Hardware,” *IEEE Access*, vol. 13, 2025.
- [15] R. Giuliano *et al.*, “Indoor Localization System Based on Bluetooth Low Energy for Museum Applications,” *Electronics*, vol. 9, no. 6, 2020.
- [16] Y. Zeng *et al.*, “Attention-Based TCN for LOS/NLOS Identification Using UWB Ranging and Angle Data,” *Electronics*, vol. 15, no. 2, 2026.
- [17] M. A. Shalaby *et al.*, “Calibration and Uncertainty Characterization for Ultra-Wideband Two-Way-Ranging Measurements,” in *Proc. of the IEEE Int. Conf. Robot. Autom. (ICRA)*, 2023.
- [18] A. Ledergerber *et al.*, “Angle of Arrival Estimation based on Channel Impulse Response Measurements,” in *Proc. of the IEEE/RSJ IROS*, 2019.
- [19] J. Kaden *et al.*, “Collecting Human Motion Data in Large and Occlusion-Prone Environments using Ultra-Wideband Localization,” in *Proc. 7th Worksh. Long-term Human Motion Prediction LHMP, IEEE ICRA*, 2025.
- [20] B. Lee *et al.*, “Comprehensive outdoor UWB dataset: Static and dynamic measurements in LOS/NLOS environments,” *Scientific Data*, 2025.
- [21] A. Kemptner *et al.*, “Automated Testbed for Repeatable Evaluation of Ultra-Wideband Localization Performance,” in *Proc. of the WFCs*, 2026.
- [22] B. Van Herbruggen *et al.*, “Single Anchor Localization by Combining UWB Angle-of-Arrival and Two-Way-Ranging: an Experimental Evaluation of the DW3000,” in *Proc. of the 14th ICL-GNSS Conf.*, 2024.
- [23] M. Naseri *et al.*, “Machine Learning-Based Angle of Arrival Estimation for Ultra-Wide Band Radios,” *IEEE Commun. Lett.*, vol. 26, no. 6, 2022.
- [24] T. Margiani *et al.*, “Angle of Arrival and Centimeter Distance Estimation on a Smart UWB Sensor Node,” *IEEE Trans. Instrum. Meas.*, 2023.
- [25] S. Bhushan *et al.*, “Single Anchor-Based Infrastructure-less Localization Performance using UWB Radios,” in *Proc. of the 20th DCOSS-IoT*, 2024.
- [26] A. Strzoda *et al.*, “Variability of BLE advertisement packets received signal strength and delivery probability in the presence of interferences,” in *Proc. of the 12th ACM Int. Symp. DIVANet*, 2022.
- [27] Monolithic Power Systems, *MA600A Digital Magnetic Angle Sensor Datasheet*, 2025, <https://www.monolithicpower.com/en/ma600a.html>.
- [28] A. Skuric *et al.*, “SimpleFOC: A Field Oriented Control (FOC) Library for Controlling Brushless Direct Current (BLDC) and Stepper Motors,” *Journal of Open Source Software*, vol. 7, no. 74, 2022.
- [29] SparkFun, “Qwiic Connect System,” <https://www.sparkfun.com/qwiic>.
- [30] H. Kv, “Notes on the Use of Propagation of Error Formulas,” *Journal Research National Bureau of Standards, Sect. C*, vol. 70, no. 4, 1966.
- [31] Nordic Semiconductor, “nRF Connect SDK (v3.2.4),” <https://github.com/nrfconnect/sdk-nrf/tree/v3.2.4>.
- [32] J. H. Choe *et al.*, “Addressing Relative Pose Impact on UWB Localization: Dataset Introduction and Analysis,” in *Proc. of the 24th ICCAS Conf.*, 2024, dataset: https://github.com/cjhhalla/RCV_uwb_dataset.
- [33] Y. Li *et al.*, “Using Wi-Fi magnetometers for indoor location and personal navigation,” in *Proc. of the 6th IPIN Conf.*, 2015.
- [34] I. Dotlic *et al.*, “Angle of arrival estimation using decawave DW1000 integrated circuits,” in *Proc. of the 14th WPNC Worksh.*, 2017.