



ULoc: Low-Power, Scalable and cm-Accurate UWB-Tag Localization and Tracking for Indoor Applications

MINGHUI ZHAO, TYLER CHANG, ADITYA ARUN, ROSHAN AYYALASOMAYAJULA, CHI ZHANG, DINESH BHARADIA, University of California, San Diego, USA

A myriad of IoT applications, ranging from tracking assets in hospitals, logistics, and construction industries to indoor tracking in large indoor spaces, demand centimeter-accurate localization that is robust to blockages from hands, furniture, or other occlusions in the environment. With this need, in the recent past, Ultra Wide Band (UWB) based localization and tracking has become popular. Its popularity is driven by its proposed high bandwidth and protocol specifically designed for localization of specialized “tags”. This high bandwidth of UWB provides a fine resolution of the time-of-travel of the signal that can be translated to the location of the tag with centimeter-grade accuracy in a controlled environment. Unfortunately, we find that high latency and high-power consumption of these time-of-travel methods are the major culprits which prevent such a system from deploying multiple tags in the environment. Thus, we developed ULoc, a scalable, low-power, and cm-accurate UWB localization and tracking system. In ULoc, we custom build a multi-antenna UWB anchor that enables azimuth and polar angle of arrival (henceforth shortened to ‘3D-AoA’) measurements, with just the reception of a single packet from the tag. By combining multiple UWB anchors, ULoc can localize the tag in 3D space. The single-packet location estimation reduces the latency of the entire system by at least 3 \times , as compared with state of art multi-packet UWB localization protocols, making UWB based localization scalable. ULoc’s design also reduces the power consumption per location estimate at the tag by 9 \times , as compared to state-of-art time-of-travel algorithms. We further develop a novel 3D-AoA based 3D localization that shows a stationary localization accuracy of 3.6 cm which is 1.8 \times better than the state-of-the-art two-way ranging (TWR) systems. We further developed a temporal tracking system that achieves a tracking accuracy of 10 cm in mobile conditions which is 4.3 \times better than the state-of-the-art TWR systems.

CCS Concepts: • **Hardware** → **PCB design and layout**; • **Human-centered computing** → *Mobile computing; Ubiquitous computing*; • **Networks** → **Location based services**.

Additional Key Words and Phrases: Ultra-Wideband, 3D-Localization, 3D-tracking, Angle of Arrival, L-shaped, Anchor Board, Hardware, cm-Accurate, Low-power, Real-time

ACM Reference Format:

Minghui Zhao, Tyler Chang, Aditya Arun, Roshan Ayyalasomayajula, Chi Zhang, Dinesh Bharadia. 2021. ULoc: Low-Power, Scalable and cm-Accurate UWB-Tag Localization and Tracking for Indoor Applications. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 3, Article 140 (September 2021), 31 pages. <https://doi.org/10.1145/3478124>

1 INTRODUCTION

Localization and tracking applications have been on an increase for various indoor settings ranging from indoor navigation [66] to asset and personnel tracking for industrial settings like hospitals [1, 11] and construction [18]. To this extent, RF-based systems [27, 46, 55, 65, 80] have been popular as their alternative, VR and camera-based systems depend on ideal lighting, are sensitivity to blockages, and are very expensive. Amongst these RF-based

Author’s address: Minghui Zhao, Tyler Chang, Aditya Arun, Roshan Ayyalasomayajula, Chi Zhang, Dinesh Bharadia, University of California, San Diego, USA, (miz107,tkchang,aarun,roshana)@ucsd.edu,zhangchi866@gmail.com,dineshb@eng.ucsd.edu.

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

© 2021 Association for Computing Machinery.

2474-9567/2021/9-ART140

<https://doi.org/10.1145/3478124>

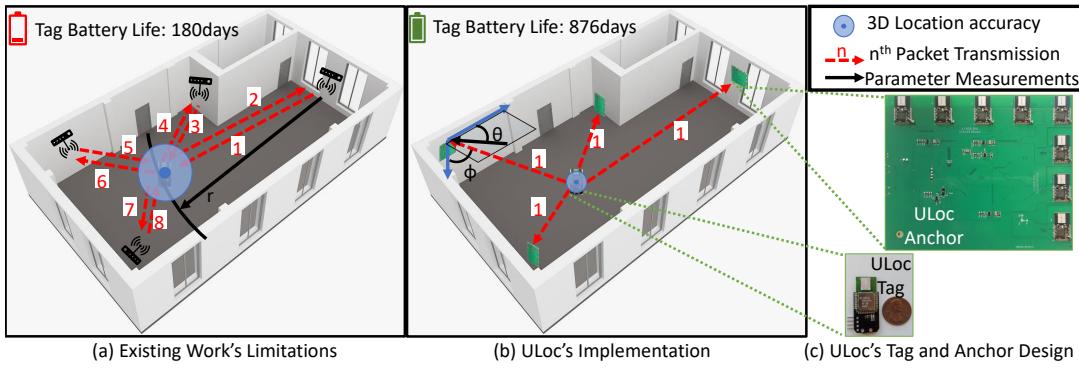


Fig. 1. ULoc: Comparing the traditional UWB localization system to ULoc, showing (a) Low-battery life (180 days) and high latency (8 packet exchanges) (b) High batter life (876 days) and low latency (1 packet exchange) (c) ULoc's custom localization anchor and tag boards, where tag is placed next to a penny, measures only 1.6×3.7 cm.

systems, Ultra-Wide Band (UWB) based systems have gained more traction in both industry [1, 8–10, 12–16, 20, 28] and academia [29, 48, 49, 105] as UWB is developed specifically as a localization-based RF protocol. This traction from industry and academia has led to the deployment of UWB based systems in smartphones [3], tags [4, 17, 19] and access points [5] to enable localization in them.

To enable these modern applications along with the localization capability in these modern devices, a new set of requirements needs to be met. We would next discuss these requirements and show where the current state-of-the-art does not meet these requirements:

R.1 **Low-latency and Real-Time Estimation:** Many of the industrial and commercial applications mentioned above require the location estimates of the UWB tags to be acquired in real-time with a location update rate of up to a few 100's of Hz [45, 95] to make time-critical and safety decisions. For example, consider the situation of an automated robotic arm in an industrial setting that loads and unloads the inventory in close vicinity of a worker. The localization system should be able to get an estimate of both the robotic arm and the worker in real-time to be able to avoid any crashes or injuries. While the traditional Two-Way Ranging (TWR) based UWB localization systems [24, 25, 29, 40, 41, 43, 46, 48, 50, 57, 58, 60, 62, 64, 71, 72, 83, 91, 100, 105] make use of measurements of Time of Flights as shown in Figure 1a and trilaterate the UWB tag. These algorithms need at least 12 individual packet exchanges across all the anchors to get the 3D location of a single UWB tag¹ limiting the real-time performance of the system and reducing the location update rate to 76Hz. Thus these traditional TWR based UWB localization systems cannot be real-time to meet the time-critical and safety requirements for these applications.

R.2 **Infrastructure-aware:** In addition, the industrial, construction and hospital settings need the infrastructure to be able to locate multiple tags for inventory and personnel analytics and management. To avoid collision in an industrial setting we talked about earlier, for example, the ‘infrastructure’ operating the machinery needs to know the position of all the equipment and personnel. While, many of the recent systems [46, 49, 61, 74, 84] reverse the localization from infrastructure-driven to ‘Tag-driven’ to improve the latency compared to the traditional TWR systems, which have very poor latency as discussed in R.1. These tag-driven systems make the tag actively transmit and receive from multiple anchors in the environment and depend on the tag to play a central role in localizing and tracking itself with assistance from the infrastructure and the anchors similar to GPS localization on a smartphone. Further, this tag-driven localization requires an additional transmission of

¹4 Anchors with 3 packet exchanges for each

each tag's location back to the infrastructure which would create more traffic and make the management of multiple tags non-real-time (R.1).

R.3 **Long operation lifetime:** Further, these tags on personnel or equipment in the industry need to be low-power to enable long-battery based operation lifetime of these tags. The tag-driven systems discussed above [46, 49, 61, 74, 84], scale well with an increasing number of tags for latency but make the tag power-hungry since all the computation is performed on the tag end to locate itself. Further, these systems require the tag to passively listen to the transmissions from multiple anchors in the environment to perform self-localization based on Time Difference of Arrival (TDoA) [46] or Angle Difference of Arrival (ADoA) [49]. This requires the tag to always actively listen for and receive incoming packets making the tag 1.8× more power hungry [6] as compared to a tag that purely transmits beacons and relies on the infrastructure for localization and tracking.

R.4 **3D Accuracy:** Finally, these UWB based systems need to be able to accurately track the UWB tags to cm-grade accuracy in 3D to avoid causing damage and crashes for example in the industrial setting we have discussed in R.1. Furthermore, in these safety-critical scenarios it is essential to be robust even in a corner case (i.e. 90th and 99th percentile error of sub-50 cm) [52, 92]. Unfortunately, most of the state-of-the-art systems [46, 49, 91] only perform localization in 2D and have had high errors for Non-Line of Sight (NLoS) cases, which cause the high errors (>sub-50cm) in these corners (90th and 99th percentile) cases.

In this paper, we present ULoc, a UWB based 3D system that achieves low-power, real-time, and cm-accurate tracking. ULoc's design makes the tag low-power (R.3) by offloading all the complexity from the tag to the infrastructure and in the process making the location available to infrastructure (R.2). ULoc achieves this by designing an anchor board that off-loads all the complexity to the infrastructure, that also enables providing the tag's cm-accurate 3D location (R.4) in real-time (R.2) with just one single transmission from the tag as shown in Figure 1b. ULoc thus satisfies the above-described requirements (R.1-R.4) by making contributions on multiple ends, including tag's protocol, anchor design, and algorithm design.

Low-Power and Low-Latency tag design: Firstly, In ULoc, we make the UWB tags low-power and low-latency by offloading all the complexity from the tag to the anchor board that we design. To do that, we enable the anchor boards in tandem with a back-end infrastructure to locate the tag accurately with a single 'Blink' broadcast transmission. This broadcast transmission is then received by all of the ULoc's anchor boards in the environment that can then enable the tag's location estimation at the infrastructure. We thus, reduce a 12 packet exchange to a single transmission, making the system real-time. Further, also make the tag low-power by just transmitting instead of receiving which is 1.8× more power-hungry, and by avoiding adding any compute requirements on the tag. Thus making the tag low-power and the localization estimation of the tag low-latency.

ULoc Anchor Design: As mentioned earlier, we enable the tag to just broadcast a single 'Blink' transmission to enable 3D accurate location estimation in real-time. We enable this with our ULoc's anchor board design that enables '3D-AoA'² estimation at each anchor that is combined across them to get accurate 3D tag location in real-time. While there have been recent systems [25, 49, 94] that have started looking into AoA based localization for UWB systems, these are still power consuming (R.3) and do not perform 3D AoA(R.4) estimations. So, the first contribution of ULoc comes from its custom-built hardware design to enable accurate 3D AoA estimation. In ULoc, we design an L-shaped 2D antenna array, that simultaneously receives the 'Blink' packet transmitted by the UWB tag across each antenna on each anchor. We further make design adjustments on our hardware to enable time, frequency, and phase synchronization across all the antennas on a single anchor board to make accurate 3D AoA estimation at each anchor board. We enable this by designing clock distribution circuits and timing synchronization circuits that overcome the limitations of the current commercial UWB transceivers [7].

²Polar and Azimuth angle measurements to enable 3D Localization, so 3D AoA. In literature, often 2D-AoA is used for just estimating azimuth angle only, as it provides 2D location. Hence, we use '3D AoA' for measuring azimuth and polar angle.

The readers can find our open-sourced hardware and firmware design updates to enable these accurate 3D AoA estimates online³.

ULoc's Algorithm Design: To enable accurate 3D AoA estimation for real-time localization, in ULoc's algorithm design we first utilize ULoc's hardware design to get the time, frequency, and phase synchronized channel impulse responses (CIR) of the 'Blink' transmissions of the tag across all the anchors. These CIR estimates across all the antennas for one single anchor board are then utilized to estimate 3D AoA at each anchor. While there are MUSIC and other high-resolution algorithms that provide accurate multipath-free 3D AoA in WiFi localization systems [55, 75, 96, 97], these algorithms need to compute heavy Singular Value Decomposition and/or averaging across multiple packets [55]. In contrast, ULoc takes advantage of the high-sampling rate enabled First Peak Index(FPI) detection system of the commercial Decawave-chipset [35]. This FPI index is accurate at least 1 nsec due to the high bandwidth of UWB systems, thus enabling multipath free 3D AoA estimation by using simple and fast 2D FFT algorithms [59].

Combining Across multiple anchors at the infrastructure: Finally, while this FPI + 2D FFT approach to estimation provides decently accurate 3D AoA estimates, these are not accurate enough to provide cm-grade localization and tracking accuracy that is required for our systems. One of the reasons for these inaccuracies in the 3D AoA estimates provided by FPI + 2D FFT is due to high mobility and sudden change in the multipath scenarios that lead to leakage of multipath signal power into the direct path tap index estimated by the FPI algorithm. The second key observation that ULoc makes to overcome this is to utilize the temporal consistency of these systems, and while Kalman Filtering [70, 93] based approaches can be readily applied, the 3D AoA errors, we have observed are non-Gaussian in nature. Thus, we develop a novel non-Gaussian temporal tracker for UWB tag location, enabling ULoc to achieve low-power and cm-accurate 3D localization in real-time.

To evaluate the efficacy of ULoc, we have deployed up to 4 of these anchors in three different 3D spaces with the ground truth obtained from VR headset [22], with random human motion experiments. First is an indoor office/home setting space with a lot of monitors, desks, and chairs spanning $10ft \times 16ft \times 10ft$, second is a large indoor open space and the third space is a larger scale indoor setting scenario with clear demarcated boundaries and few heavy reflectors occurring frequently in a space of $20ft \times 20ft \times 13ft$. We have located and tracked our custom-designed small-form-factor UWB tag (ULoc's tag), with custom firmware using the same DW1000 UWB transceiver, under multiple scenarios to achieve the following results:

- Comparing with the state-of-the-art two-way ranging-based localization technique (TWR), ULoc achieves a median (90th percentile) 3D localization error that's $1.8 \times$ ($2.5 \times$) smaller in stationary conditions. Furthermore, in high mobility conditions, ULoc achieves a median (90th percentile) error $3.1 \times$ ($4.3 \times$) smaller than that achieved by TWR. This satisfies R.4.
- ULoc also achieves a localization latency of 1 msec per tag location, which is $13 \times$ lower than the state of the art two-way ranging based techniques (TWR) that can at the best give a localization latency of 13 msec. Thus showing the scalability of ULoc. This satisfies R.1.
- ULoc can achieve this localization latency and accuracy with minimal power consumption wherein, ULoc's algorithm consumes $31 \mu J$ of energy for one location estimate and the state of the art two-way ranging based trilateration algorithm (TWR) consumes up to $286 \mu J$ of energy per location estimate. This satisfies R.3.
- ULoc can track multiple tags in real-time without any loss in its accuracy. We show this via simulating a simple medium access control (MAC) protocol and testing 4 co-located tags in an office environment. This satisfies R.2.

³<https://github.com/ucsdwcsng/ULoc-public>

With these results, we show that ULoc makes the following key contributions to UWB-based 3D localization. First, we expand on existing 2D-FFT-based algorithms for polar and azimuth angle estimation to provide few-centimeter accurate tracking even in high-mobility conditions in Section 2. Second, we develop multiple UWB transceiver designs that enable accurate angular estimation in both azimuth and polar in Section 3. Further, we robustly test ULoc’s algorithm and platform in a variety of mobility and environment conditions in Section 4. Finally, we present our intended future work to tackle some of the existing limitations in Section 6.

2 ULOC’S ALGORITHM AND PROTOCOL

In this section, we present the details of ULoc’s protocol and algorithm design that enables ULoc to be low-latency, low-power and cm-accurate 3D localization. As described in section 1, one of the ULoc’s contribution that enables these features is its novel anchor board design which is described in section 3. This hardware design in tandem with the tag’s protocol of single ‘Blink’ transmission that enables infrastructure-aware tag localization by transferring all the complexity from the tag to the anchor. We first detail this protocol of the tag that enables this low-power and low-latency localization in section 2.1. After we present the protocol design on the tag that enables low-power and low-latency localization, we then do a deep dive into our algorithm design of the computations at each individual anchors that enables real-time 3D AoA estimation in section 2.2. Finally, we present how the infrastructure combines these real-time 3D AoA estimates from multiple of the ULoc’s anchor boards to achieve cm-accurate 3D location at the infrastructure for all the tags in the environment in section 2.3.

2.1 Enabling Low-power and Low-latency Tag Design and Protocol

In ULoc, we design our UWB tag to be low-power and enable low-latency location updates. To enable the above, we would first summarize the trade-off in the current UWB based localization systems. Firstly, UWB transceivers consume $1.8\times$ more power while receiving than while transmitting [6], clearly showing that the tag should be transmitting more than receiving. A more detailed description and comparison of these power-consumption metrics can be found in appendix B. Furthermore, traditional TWR systems [31, 32, 46, 49, 72, 91, 103] employ at least 12 packet transfer including running the tag in receiver mode continuously to perform localization making the tag, power consuming and the system, high-latency. While, recent systems [31, 32, 46, 49, 72, 91, 103] are ‘tag-driven’ to ensure low-latency where only the tag knows its location, making the tags require compute and thus power-hungry. Further the tag has to receive simultaneously from multiple anchors and should constantly be on listening mode, making it hard for it duty cycle its sleep mode and also as mentioned earlier that receiving consumes more power. Thus, we see a trade-off of low-power for low-latency in the current UWB localization systems by enabling infrastructure-aware and infrastructure-driven tag localization in contrast to the more power-hungry ‘tag-driven system’.

Taking this trade-off into account, we design a low-power and low-latency system by pushing the complexity of location computation and scheduled packet transfer protocols onto the Anchor or the infrastructure of our system. We enable this by enabling a localization protocol, wherein the tag just needs to transmit a one-single packet to let the infrastructure figure out the location of the tag. These single packet tag transmissions are received at multiple ULoc anchors spread across the environment, which compute the azimuth and polar angle-of-arrival of the signal (3D-AoA). The 3D-AoA estimates across multiple anchors are then used to locate the tag in real-time.

Furthermore, we can enable blink without modifying the UWB protocol as follows, in ULoc’s protocol, the first time the tag is booted, it initiates standard protocol compliant repeated transmissions (*Blinks*) to inform the infrastructure of its existence. Upon discovery of the tag, the anchors transmit *Range-Init* response as is common with the UWB protocol. During the *Range-Init* transmission, the anchors program the tag to only transmit a *Blink* at a periodic interval to enable accurate localization of the tag. Thus, the tag only needs to transmit only one *Blink* message to let all the ULoc anchors receive it and locate it simultaneously. This removes compute from the

tag reducing the power-consumption on the tag. Additionally, this *Blink* interval can be modified based on the number of co-located tags in the environment to avoid packet collision. Finally, as mentioned earlier, the 3D AoA and the tag location estimations are performed at the infrastructure removing all complexity from the tag.

Thus the single ‘Blink’ transmission-based protocol makes the tag last longer up to 2.4 years, compared to the state of the art tag driven systems [31, 32, 46, 49, 72, 91, 103] that only last as long as 4 months, or the traditional TWR based systems [24, 40, 41, 60] that last only 3.1 months. Further, this protocol also makes the location update rate real-time for ULoc with a latency of only 1msec for each location for a tag, where the traditional UWB systems have a latency as high as up to 13msec.

2.2 ULoc’s Anchors Algorithm: Real-time 3D AoA Estimation

In ULoc, we thus use the above-described tag protocol in tandem with ULoc’s novel 8-antenna L-shaped anchor board design shown in Figure 1b to estimate 3D AoAs by making simple and smart design choices that make our final localization real-time, low-power, and cm-accurate. A detailed description of ULoc’s hardware is described in section 3, ULoc’s anchor boards shown in Figure 1b enable accurate time, frequency, and phase synchronization necessary for accurate 3D AoA estimation [55, 59, 97].

In this section, we first tackle the issue of making the 3D AoA estimates in real-time using our anchor board design. ULoc’s AoA estimation design is inspired by the decade long research in WiFi-based systems that use algorithms such as MUSIC [76, 97], SpotFi [55], ESPRIT [75], and FFT inspired algorithms [26, 59, 96] that estimate the 2D AoA of the incoming signal at the receiver. In a similar vein, authors [59] extend to 3D AoA estimates by employing a two-dimensional antenna array at the receiver. Unlike UWB systems, WiFi-based systems struggle with limited bandwidth (20-80MHz) that make multipath resolution challenging for WiFi-based localization systems. More specifically due to limited bandwidth, the WiFi-based systems can resolve direct-path from multipath only when they are separated by more than 3.75m.

To tackle this, authors [55] have used novel CSI-smoothing and 2-D MUSIC algorithms to achieve multipath resolution. More specifically, MUSIC-based algorithms estimate auto-correlation and then use the SVD to define the noise space. The signal-projection orthogonal to this noise space leads to the parameter (like 3D-AoA) estimate [55, 97]. Further, most of these MUSIC algorithms perform combining of data across space/antennas [97] or across time/packets [55] to achieve decimeter level localization performance.

In ULoc, we want to minimize the latency and so firstly we would not want to average across multiple packets to compute the auto-correlation as that increases the latency for single location computation. Unfortunately though, without averaging and using only one packet leads to noisy estimates for auto-correlation and thereof 3D AoA values. We further present the same in our evaluation Figure 10c. Secondly, the SVD computation involved in MUSIC requires an $O(n^2)$ floating-point operations, whereas the FFT-based algorithms [26, 59] (if the direct path can be identified) requires only $O(n \log(n))$, where n is the number of input data points. Based on these observations we choose to rely on FFT-based algorithms for ULoc’s 3D AoA estimation.

Above, we assume the FFT-based 3D AoA estimation that we can extract the direct path’s tap across all the antennas from the given CIR information from each antenna of our anchor board. To get the direct path’s tap across all the antennas we rely on the design of the antenna chipset [7]. These antenna chipsets provide the First Path Index (FPI) which is the direct path’s index as the direct path travels the least amount of time. This FPI estimate provided by the antenna chipsets has a resolution of up to 1 nanosecond. In Figure 2(a), we show the channel magnitude measured across four antennas at ULoc’s anchor in the absence of any multipath. We compute the likelihood profile (shown next to it) of the azimuth and elevation angle of arrival at this time-index (FPI) and observe that the 3D-AoA estimate (blue circle) matches closely with ground truth (purple triangle). Thus, FPI provides a reasonable estimate of the direct path and therefore it could be used with FFT-based 3D AoA estimation. Based on these observations we extend the FFT-based 3D AoA estimation presented in [59] to

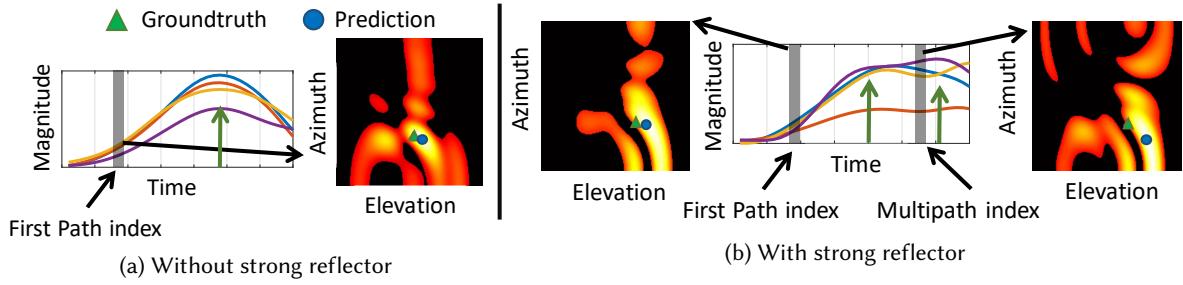


Fig. 2. Effects of a strong reflector placed 2 feet away from ULoc's anchor in Env-3 as shown in Figure 7. (a) Shows the CIR estimates and 3D AoA profile without the reflector where we see one clear peak. (b) Shows CIR and 3D AoA profile after we add the strong reflector. The green-arrow shows the incoming direct path and multipath signal. We observe that the FPI's fine-time resolution allows ULoc to robustly resolve for multipath in the environment.

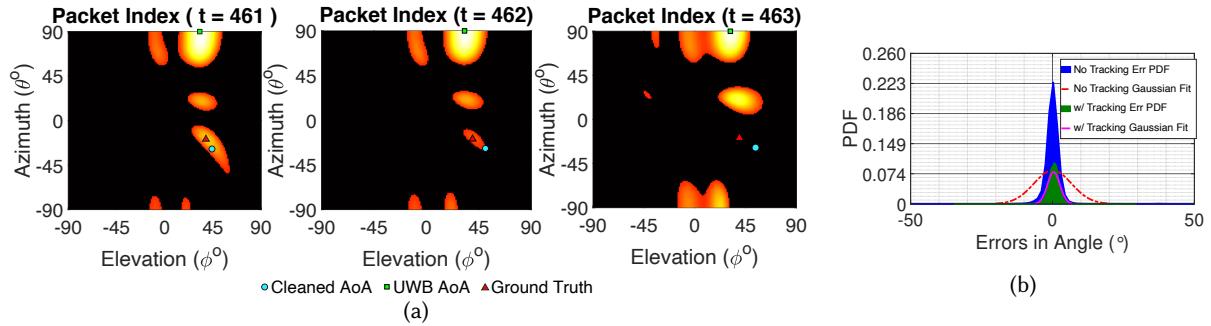


Fig. 3. (a) Intuition: ($t = 461$): Shows likelihood map with existent correct peak indicated by red triangle. ($t = 462$): Shows likelihood map with existent, but fading correct peak. ($t = 463$): Shows likelihood map with correct peak non-existent. (b) AoA error distributions: Shows the error distributions and their corresponding Normal fit performed on errors for UWB based measured angle of arrivals before applying ULoc's AoA tracking algorithm (called UWB), we can see that the errors in angle of arrival is very weak fit because of high errors at the tails. And after applying ULoc's AoA tracking algorithm (called ULoc) we can see that the AoA errors are normally and tightly distributed.

FPI + 2D FFT-based 3D AoA estimates in real-time. We present a more formal mathematical formulation and the axis representation we use to achieve our FPI + 2D-FFT based 3D AoA estimates for our L-shaped anchor board design in Appendix Section A.

2.3 Accurate, Real-time 3D AoA Estimation

The above 3D AoA estimates ($\theta^{\text{UWB}}, \phi^{\text{UWB}}$) are provided in real-time and are more accurate than MUSIC-based estimates without averaging. Unfortunately, we observe in a few 90th percentile data samples that when multi-path gets closer it disintegrates into two theta-phi values and our initial assumption of FPI capturing the direct path is invalid. To understand this better we show few examples of these 90th percentile cases with close multipath. We expect that each consecutive antenna would see a constant phase difference for the vertical or horizontal antenna array. If we have these constant phase differences, the theta-phi likelihood-map would have a single peak. This single peak provides a confident estimate of the azimuth and elevation angle-of-arrival. Unfortunately, in certain scenarios, close existing multipath in the channel can corrupt the phase difference between some pairs of antennas. This will result in multiple ambiguous peaks in the theta-phi likelihood-map shown in the profile in Figure 3(a). In a few more instances the peak corresponding to the actual ground truth completely vanishes as

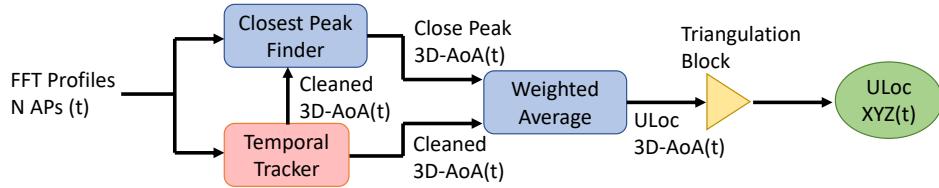


Fig. 4. Overall Algorithm Flow: Given the initial input likelihood-maps across N-anchors from Eq 5, separately determine 3D AoA and triangulate and find ULoc localization. Lower block gives an in depth view of how the temporal tracker takes input likelihood-maps from Eq 5 to find a cleaned 3D AoA. We refer to the anchors as "APs" in this figure.

shown in Figure 3(a)($t = 463$). These erroneous peaks can lead to incorrect 3D AoA predictions, particularly in high mobility situations where the peak can move greatly throughout the likelihood-map within a single time instance.

To overcome this phase distortion, we leverage the continuous nature of human or object motion in the environment and choose to utilize this to our advantage to track the 3D AoA predictions over time. While a simple solution would be to pass the 3D localization estimates predicted using the 3D AoA based trilateration estimates from multiple of ULoc's anchor boards through a simple Kalman Filter [70, 93]. While these algorithms try to model the errors in the 3D localization to track the user accurately over time, they do not accurately model the underlying error distributions of the 3D AoA systems. One of the major reasons for these errors is that these 3D AoA errors are non-Gaussian as shown in Figure 3(b).

The non-Gaussian error distribution of AoA based systems can be attributed to the fact that AoA estimation algorithms equate the phase difference to $\sin()$ and $\cos()$ of AoA, therefore resulting in higher error for θ and ϕ at angles close to $+/ - 90^\circ$. This non-uniformity of AoA results in high errors at angles beyond $\pm 60^\circ$ leading to a non-Gaussian-like spread towards the higher end of the angular spectrum. Hence, we observe an over-estimation in the error of the measurements and a poorer fit of the Kalman filtered path.

A brief design of the novel temporal tracking algorithm employed in ULoc is depicted in Figure 4. The first and foremost intuition is within the temporal tracker that tracks the user over-time by estimating the user's velocity from the previous $N=7$ estimates, and uses that to predict the user's next location, based on the optimized user location from the previous temporal estimate. In our temporal tracker, we utilize the global maxima peak for the previous N time instances rather than our optimized prediction from the previous instances. We take this approach to overcome the case that no good peak exists for an extended period of time, for when using our previously optimized peak would rely mostly on the low inertial cleaned prediction. This could result in our optimized prediction drifting off forever in the direction it was last moving. For example, consider the peaks shown in Figure 5(c), where the correct peak disappears. This disappearance of the peak over multiple packets with the optimized prediction would continue to drift on the same course with the same velocity as the last measurement.

To define the temporal tracker more formally, let us consider the maxima peaks derived from our algorithm in section 2.2 at time instances $(t-n), n = 1, 2, \dots, N=7$ for all the N_{ap} anchor boards in the environment, $(\theta_{t-n}^{\text{UWB}}, \phi_{t-n}^{\text{UWB}})$. We then estimate our Maxima 3D XYZ positions by simple trilateration [55, 97] ($XYZ(t)$), which are then utilized to get the estimate of the approximate velocity of the UWB tag $v_{\text{temporal}}(t)$. Using this velocity along with the previous instant's predicted location, $XYZ^{\text{temporal}}(t-1)$ the temporal tracker predicts the location of the current instant $XYZ^{\text{temporal}}(t)$ as $XYZ^{\text{temporal}}(t-1) + v^{\text{temporal}}(t) * \Delta t$, which is then remapped to the 3D AoA estimates for all the N_{ap} anchors in the environment as $(\theta_t^{\text{cleaned}}, \phi_t^{\text{cleaned}})$.

Depending on the temporal tracker alone can be harmful as we can completely neglect the current 3D AoA estimation profile, which would lead to incorrect predictions during sudden changes in direction and orientation

of the tag motion. Thus in addition to the temporal tracker, we use the prediction of the temporal tracker to estimate the closest peak to it amongst the peaks in the current profile in the ‘Closest Peak Finder’ block shown in Figure 4. We perform this operation to weigh our trust over the current profile as well, and not rely completely on the temporal tracker alone. We then take a weighted average of both the temporal tracker 3D AoA estimate ($\theta_t^{\text{cleaned}}, \phi_t^{\text{cleaned}}$) and the closest peak finder 3D AoA estimate ($\theta_t^{\text{close}}, \phi_t^{\text{close}}$) to get accurate real-time 3D AoA estimate of ULoc, ($\theta_t^{\text{ULoc}}, \phi_t^{\text{ULoc}}$). These weights, $W_{\text{reliability}} (< 1)$ define the reliability of the current 3D AoA estimates of our temporal tracker ($\theta_t^{\text{cleaned}}, \phi_t^{\text{cleaned}}$).

Weight Generation: The weighting factor ($W_{\text{reliability}}$) describes the error distribution of each likelihood-map, using the weights (peaks, angular, and distance) to account for non-Gaussian 3D-AoA errors. The weights themselves ($w_{p,t}$, $w_{a,t}$, and $w_{d,t}$) are averaged to determine the weighting factor. The three weights take in three inputs: theta cleaned, theta close, and the number of peaks in the likelihood-map at time t. The first weight ($w_{p,t}$) counts the number of peaks in the likelihood-map, and grants less reliance on theta close as the number of peaks increase. The second weight ($w_{a,t}$) accounts for erroneous 3D-AoA estimation following Eq 5 close to $\pm 60^\circ$, granting less reliance to theta close as the average estimation between theta cleaned and theta close approaches 3D-AoA extrema. The third weight ($w_{d,t}$) compares the difference between theta cleaned and theta close, granting less reliance to theta close as the difference increases. Given the high inertial behavior of theta cleaned, this weight accounts for large instantaneous changes in the likelihood-map, particularly in disappearances of the correct peak.

Re-Initialization of Temporal Tracker: A final consideration in determining accurate 3D-AoA is re-initialization of theta cleaned ($\theta_t^{\text{cleaned}}$). Given the differential determination of theta cleaned, drift may occur and cause highly erroneous 3D-AoA estimation. It thus becomes necessary to re-initialize our theta cleaned estimation alongside reliable theta close (θ_t^{close}) estimations. Theta close is considered reliable when within $\pm 60^\circ$, and the weighting factor is within 0.35 ($W_{\text{reliability}} <= 0.35$). Under these conditions, we set theta cleaned to be the same as theta close, $\theta_t^{\text{cleaned}} = \theta_t^{\text{close}}$.

ULoc 3D-AoA: ULoc’s final 3D AoA determination is given by Eq 1 below.

$$\theta_{k,t}^{\text{ULoc}} = W_{\text{reliability}} \theta_{k,t}^{\text{cleaned}} + (1 - W_{\text{reliability}}) \theta_{k,t}^{\text{close}}; \quad \phi_{k,t}^{\text{ULoc}} = W_{\text{reliability}} \phi_{k,t}^{\text{cleaned}} + (1 - W_{\text{reliability}}) \phi_{k,t}^{\text{close}} \quad (1)$$

$$\text{where, } W_{\text{reliability}} = \frac{(w_{d,t} + w_{p,t} + w_{a,t})}{3} \in [0, 1]$$

where $\theta_{k,t}^{\text{ULoc}}$ is the final theta estimation for anchor $k \in \{1, 2, \dots, N_{\text{anchor}}\}$ at time index, t . Thus ULoc achieves accurate Real-time 3D AoA estimates from a single anchor board

2.4 Accurate, Real-time 3D Localization

For a multi-anchor setup ULoc’s anchors do not need to perform any TWR measurements, & simply employ a standard triangulation algorithm using the accurately tracked 3D AoA estimates ($\hat{\theta}_k, \hat{\phi}_k$) from multiple anchor points ($\forall k = 1, 2, \dots, N_{\text{anchor}}$), where we define the line l_k starting from the anchor’s location, \mathbf{x}_k in the direction of ($\hat{\theta}_k, \hat{\phi}_k$) as

$$l_k \equiv \mathbf{p}_k = \mathbf{x}_k + t\vec{v}_k \quad (2)$$

where, \mathbf{p}_k is any point along the line l_k for a given translation of t away from the known point on the line, \mathbf{x}_k , and $\vec{v}_k = [\sin(\hat{\phi}_k), \cos(\hat{\phi}_k) \sin(\hat{\theta}_k), \cos(\hat{\phi}_k) \cos(\hat{\theta}_k)]$ is the unit vector defined in the direction defined by ($\hat{\theta}_k, \hat{\phi}_k$). With this definition of the line, l_k , we can estimate the location of the tag \mathbf{x}_{tag} as

$$\mathbf{x}_{tag} = \operatorname{argmin}_{\mathbf{x}} \sum_{k=1}^{N_{ap}} (\mathbf{x}_k - \mathbf{x})^T (I - \vec{v}_k \vec{v}_k^T) (\mathbf{x}_k - \mathbf{x}) \quad (3)$$

where I is an 3×3 identity matrix. Using this formulation we can perform much faster accurate 3D localization that is at least $20\times$ faster than a standard TWR based localization algorithm.

3 ULOC'S HARDWARE, SOFTWARE DESIGN AND SOFTWARE PLATFORM

In the previous section, we assumed that ULoc's platform can provide anchors, which can measure the 3D angle-of-arrival for any UWB transmission. In this section, we would present details of ULoc hardware design, firmware design and tag design, wireless calibration and data processing to enable a real-time operation. ULoc platform would be open-sourced allowing the community to build large scale anchors to be forefront of this exciting direction in indoor 3D localization.

3.1 Anchor Hardware Design

The goal of ULoc anchor is to achieve 3D AoA by receiving single packet from the tag. To measure accurate 3D AoA, we need to measure the relative phase difference in the channel impulse response between all the antennas accurately. It therefore requires all the antennas which are connected to the transceivers on a ULoc anchor to be synchronized in time, frequency, and phase. However, note for ULoc to work, the multiple ULoc anchors deployed in environment don't need to be synchronized, only the transceivers on each anchor needs to be synchronized. In this subsection we detail ULoc's anchor hardware designs to achieve these goals. The open-source ULoc anchor hardware design would enable the community to quickly develop and experiment on synchronized multi-transceiver UWB anchors.

L-shaped Antenna Array: Each ULoc anchor contains 2 orthogonal arrays of commercial, off-the-shelf DWM1000 UWB transceivers, with 4 modules in the horizontal direction and 4 modules in the vertical direction as shown in Figure 5(b) to measure theta-phi AoA. Additionally, the anchor has a micro-controller, a clock distribution network/synchronization circuit, and power supply circuits. The anchor measures 16.0×13.4 centimeters. The DWM1000 module contains the DW1000 UWB transceiver chip and a ceramic antenna, along with supporting circuits [7]. One of the most important factors to consider while designing the L-shaped antenna array is to factor in the polarization of antennas in both linear arrays. Although these two linear arrays are perpendicular, we need to make sure that all the antennas are oriented the same way. Given that the form factor of each DWM1000 module is 1.3×2.3 cm, we place the modules in the same orientations such that all adjacent antennas are 3.3 cm apart, which is the half wavelength for the 4.4928-GHz channel in ULoc system, which provide distance between the antennas less than half a wavelength necessary to achieve unique theta-phi estimates.

Coordination of Transceivers in the anchor: In a normal UWB ranging system, each UWB transceiver is controlled by an individual microcontroller to handle the protocol and timing. A straight-forward way to build a multi-transceiver anchor is to have multiple such transceiver-microcontroller pairs, and then a central controller to command and time them to work together. However, this quickly makes the system expensive and complex with multiple asynchronously-running microcontrollers making timing extremely tricky. Instead, as illustrated in Fig. 5a, ULoc uses one single STM32F446RCT6 ARM Cortex-M4 microcontroller (MCU) [82] to control all of the DW1000 UWB transceivers.

The DW1000 chips and the MCU communicate via two 20 MHz Serial Peripheral Interface (SPI) buses, which is the highest speed supported by the chip, with 4 chips on each bus. The MCU selects the chips to communicate on the SPI buses in a round-robin manner, using 8 of its general purpose input/output (GPIO) pins connected to respective chip's *chip select* (CS) pin. Therefore, to communicate with a specific transceiver, the MCU asserts the GPIO pin connected to the respective transceiver's CS pin, and then starts communication on the SPI bus. In this

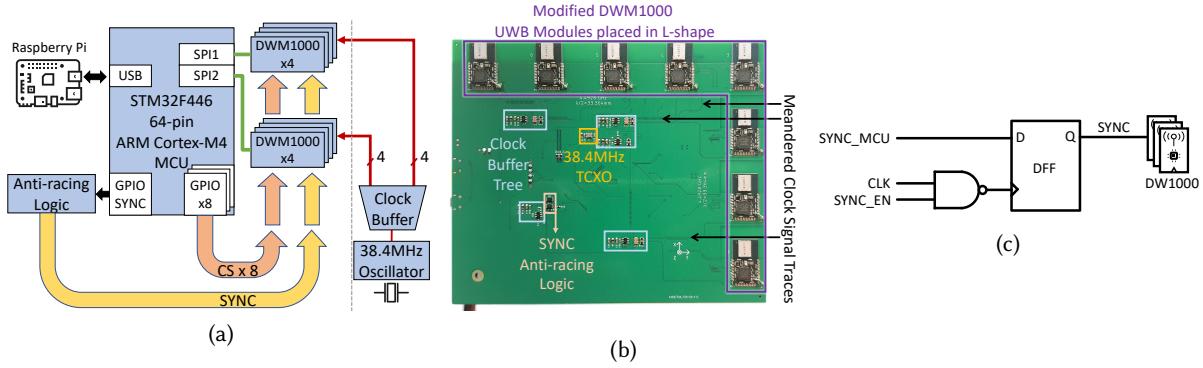


Fig. 5. ULoc Anchor Board: (a) Shows system design for anchor PCB. (b) Shows picture of anchor PCB bottom side with L-shaped antenna array and clock distribution network. (c) Shows ULoc’s anti-racing SYNC circuit logic design

setup, all the transceivers are made to work together and the MCU can handle the communication protocol as if they were one unified transceiver.

Furthermore, when Tx / Rx event occurs at a transceiver, it asserts its *interrupt request* (IRQ) pin to notify the MCU of the event so that the MCU can take action accordingly. Due to the limited amount of GPIOs on MCU and considering the scalability of using a different number of transceivers on an ULoc anchor, we connected the IRQ pins of all 4 transceivers on each SPI bus to an OR gate before connecting to a MCU input pin.

Clock Distribution and Synchronization: Under ideal circumstances, AoA should be derived from transceivers that share the same carrier clock, so that frequency and phase reference are consistent across all transceivers. However, in practice, each UWB transceiver chip generates a carrier clock with its own phase-locked loop (PLL), all of which runs independently with slightly different frequency and phase offset. Fortunately, since the PLL on each transceiver eventually aligns the carrier clock phase to the phase of the reference clock, we can synchronize the carrier frequency and phase by feeding the same reference clock to all the transceivers. The quality of synchronization depends on the quality of the reference clock signal; hence we use clock buffers and a carefully designed distribution network to avoid long clock traces running across the board, giving the reference clock signal the best quality we can achieve.

The clock on ULoc anchor is generated with an Abraccon ASTXR-12-38.4MHz temperature compensated crystal oscillator (TCXO) [23], whose output is level-shifted to 3.3 V, and then buffered/fanned out with Microchip PL133-37 clock buffers [68], before being fed into each transceiver. All clock traces are measured and meandered if necessary, as shown in Fig. 5b, to ensure the same distance that the clock signal travels from the TCXO to each of the transceivers. Note that each DWM1000 module comes with a crystal preinstalled as its clock source, so to feed all the modules with the same clock, we removed the metal shields on the modules and removed the preinstalled crystal. Then to connect each DW1000’s clock input to the clock distribution network, we solder a wire from a pad connected to the clock buffer on anchor PCB, to a pad on the module where the removed crystal’s output is.

The UWB transceiver provides an additional SYNC signal input, which can be used to reset the timestamp in each transceiver to keep them aligned. The SYNC signal is sampled at the rising edge of the reference clock signal [6]. However, without special prevention, SYNC signal can sometimes become asserted when the reference clock is close to a rising edge, and due to inevitable subtle differences in the delays in the reference clock distribution network, different transceivers could sample the SYNC signals at different reference clock cycles, which introduces an offset in the timestamps. Such an offset might be compensable in software as it can only be one reference

clock cycle, but there is a simpler way to eliminate it from the source: we can use the logic shown in Figure 5(c) to ensure *SYNC* is asserted only during the falling edge of the reference clock, giving it maximum clearance from rising edges of the clock. We chose this solution due to its simplicity and reliability.

Therefore, ULoc anchor achieves all the required time, frequency, and phase synchronization via 1) Driving the transceivers from the same clock source; 2) performing wired timebase synchronization on PCB.

PCB Layout: To ensure the quality of the clock signal, ULoc anchor is manufactured on a 4-layer controlled impedance PCB, with all critical elements including the DWM1000 modules, 38.4-MHz oscillator, clock distribution network, and synchronization circuits placed on the same side of the PCB. Furthermore, the clock trace impedance is calculated based off the trace width and PCB material, and proper termination resistors were added in series next to the clock buffer 20Ω output pin, so that the clock buffer output impedance matches the trace impedance, to avoid unwanted clock signal reflections [68].

Power Distribution: The anchor board also contains multiple 3.3-V low-dropout voltage regulators (LDO) to supply the MCU and the UWB transceivers, which are fed by an external DC power supply. Three AP1117-3.3V 1A LDOs [36] are used to power MCU, vertical 4-transceiver array, and horizontal 4-transceiver array. Multiple AP2112-3.3V 600mA LDOs [37] are used in proximity of each clock buffer chip.

The anchor has a high peak current consumption that produces sharp current spikes at more than 1A, due to the fact that all 8 transceivers receive the packet from tag at the same time. To mitigate voltage dropping in supply rails due to the current spike, we place 6-V power supplies close to each anchor, so that the supply voltage would not go below the minimum input voltage of the LDO (4.7 V), and result in transceiver power instabilities causing issues such as frequent receiver errors. Moreover, careful power decoupling is also added on the PCB in proximity to all the power inputs. The decoupling circuit involves appropriate parallel decoupling capacitors depending on the IC requirement, and bead in series, to mitigate power rail high frequency noise.

3.2 Anchor Firmware Design

The anchor firmware is based on Decawave’s firmware for TREK1000 development kit [33], which provides a demo TWR anchor and tag implementation, with one transceiver-MCU pair. As specified in the firmware package’s source code documentation, the firmware implements a state machine handling different transceiver states including setting up transceivers, sleeping, setting up Tx parameters, enabling Rx, receiving Rx data, etc.

Coordinating Multiple Transceivers: As discussed in section 3.1-*Transceiver Coordination*, ULoc anchor connects 8 transceiver modules to one single MCU via two SPI buses, and the 4 modules on each SPI bus has their IRQ pins passed through an OR gate to an MCU input port. To enable the MCU to control all transceivers in a synchronous fashion, we modified the firmware to stores states for each transceiver separately in memory. To achieve this, we create an array of *DW1000 Instances* – a data structure that stores the state information, configuration, and all data for one transceiver. Therefore, the MCU can traverse this array to access state information and communicate with the respective transceiver.

Further details on exact implementation of extracting data and outputting it to the infrastructure are detailed in appendix C.1. Using this pipeline we obtain the FPI index and FPI-1 to FPI+5 to perform raised cosine based upsampling for better resolution around the direct path’s tap. These 7 CIR taps are then received by individual Raspberry Pi’s [73], which communicate with a central server that computes ULoc’s location estimates.

3.3 Tag Design

ULoc can work with any UWB tag. We also design a low form factor tag to allow us to attach to any sensor or any object as necessary. The ULoc tag design is very simple and resembles a conventional UWB tag since the anchor handles all the complexity. In ULoc, the tag only transmits standard UWB *blink* message periodically, after the standard discovery phase. The periodicity is programmed by the anchor/infrastructure network.

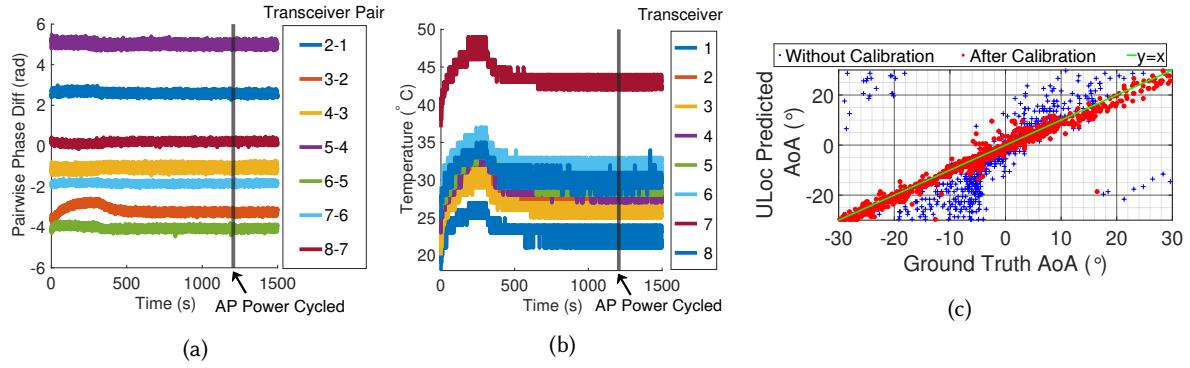


Fig. 6. (a) Shows phase differences measured across a period of 25 minutes with the tag kept stationary at zero AoA, with anchor (AP) disconnected from power and connect again at 20 minutes mark. (b) Shows the temperatures measured at the individual antennas on board in the same measurement as (a). The phase difference remains stable after restarting the host software (black vertical line), given that the temperature stays constant. (c) Shows the erroneous AoA measurements before wireless calibration and how 0-AoA based wireless calibration has stabilized our AoA measurements.

Tag Hardware ULoc tag has an STM32F072CBT6 ARM Cortex-M0 MCU [81] that interfaces with a single unmodified DWM1000 module. For simplicity, the tag receives power via a power bank through the micro USB connector. In the future, the tag can have a battery attached directly to it, and IMU/barometer sensors could be added as well. The IMU sensor could enable accurate real-time tracking of the tag. The whole tag measures 1.6×3.7 centimeters as shown in Figure 1c.

Multi-Tag Firmware The tag firmware is also based on Decawave’s TWR firmware for TREK1000 development kit. We stripped down the code to have a basic functionality, i.e. support for discovery, poll, and *blink* packets at fixed intervals, while each anchor passively listens for them with its 8 antennas. The tag’s unique hardware ID is embedded into the *blink* packet so the anchor can correctly identify different tags when multiples of which are present.

3.4 Wireless Calibration

To verify whether the on-board synchronization works perfectly, we have placed a tag at a zero theta-phi AoA relative to an anchor, and collected CIR packets for 1 hour. Ideally, the phase difference from different transceivers should be 0 when theta-phi AoA are both zero (section 2.2) after the synchronizations detailed in section 3.1. The results for first 25 minutes⁴ is shown in Fig. 6(a,b). We can see that 1) due to inevitable differences between the transceivers, a carrier phase offset exist which results in non-zero phase difference observed between adjacent transceivers at zero AoA; 2) while the phase differences drift in the first 8 minutes of the first bootup of the board, as temperature rises, the phase and temperature start to settle simultaneously afterwards; 3) the offset is not affected by power cycling anchor boards which was performed on the 20 minute mark.

Therefore, ULoc must perform a one-time wireless calibration to measure such phase offsets, so that it could be subtracted from all measurements even after reboot, which to best our knowledge has not been demonstrated. The calibration process takes around 10 minutes considering the temperature stabilization period, and any actual location measurement should also be done after the anchor’s temperature stabilizes for best performance. In the future, temperature can be sampled with the on-chip temperature sensors in DWM1000 and MCU STM32F446RCT6

⁴The results appear to be the same flat line after the initial 10 minutes

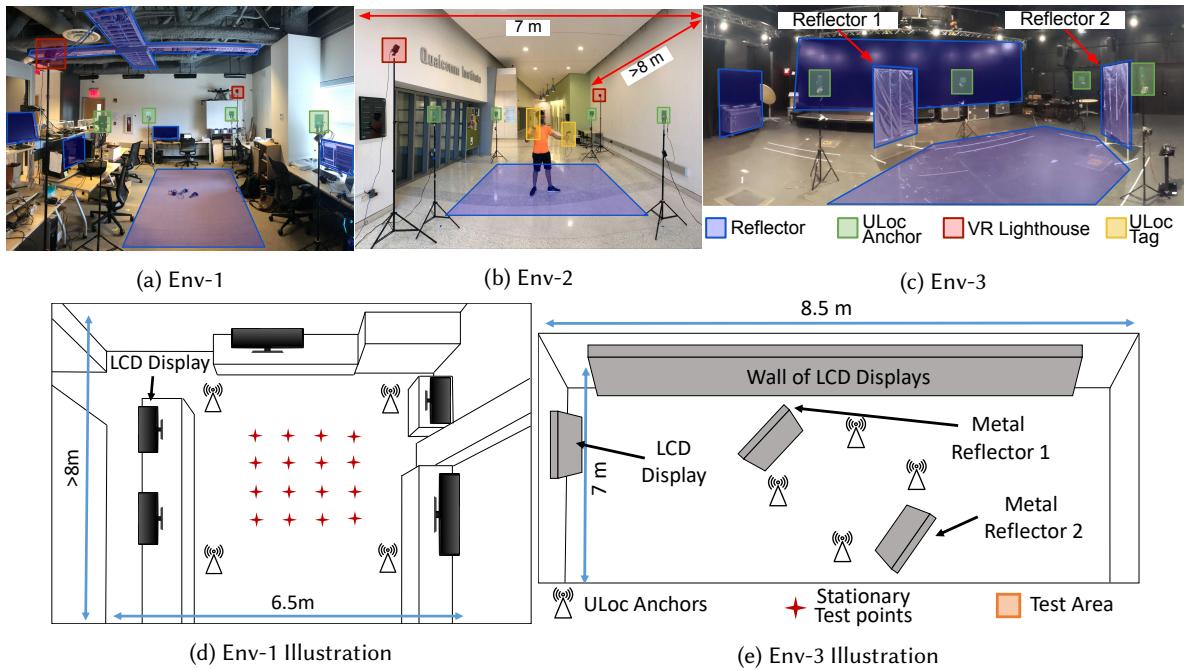


Fig. 7. ULoc was evaluated in three spaces referred to as (a, d) Env-1 (b) Env-2 (c, e) Env-3, each with varying amount of complexity, reflectors and naturally occurring environments. The blue-shaded regions are explicit large reflectors in the environment. The green-shaded regions are the ULoc Anchors. The yellow shaded regions are the UWB Tags. The red stars in (d) refer to the points where stationary data collection was performed. In all other cases, the tag is always held by a user who is making random motions with varying speed.

to allow automatic correction of such offsets even in environments with unstable temperatures. Fig. 6(c) shows the effect of the zero AoA based wireless calibration.

In summary, the 7 CIR taps from each anchor are received and are upsampled using a raised cosine filter by 8 times, which in-lieu with FPI at 15.65 FPI can then be used to extract a more accurate direct path tap. We then perform the SFD and wireless AoA calibration as discussed in section 3.4 to these 8 CIR direct path values from the 8 antennas of one anchor across all the anchors in the environment. This CIR is then utilized to perform 3D AoA estimation, and final 3D localization as described in section 2

4 EVALUATION

ULoc's system design makes contributions in both hardware and algorithm design. ULoc's algorithm design enables triangulation based accurate 3D localization and tracking for real-time systems. We evaluate ULoc in three different real-world environments and discuss ULoc's localization and tracking results in Sections 4.1- 4.3. We show that ULoc outperforms state of the art median tracking error by at least 3 \times and achieves accurate 3D tracking with a median RMS error of 8cm in 3D, and outperforms state of the art median stationary condition error by 1.8 \times , with a median error of 3.6cm. We then evaluate the latency and power consumption of ULoc compared to the state of the art UWB based localization algorithms in section 4.4 and section 4.5 respectively. Additionally, ULoc's hardware design makes 3D AoA measurements possible from commercially off the shelf

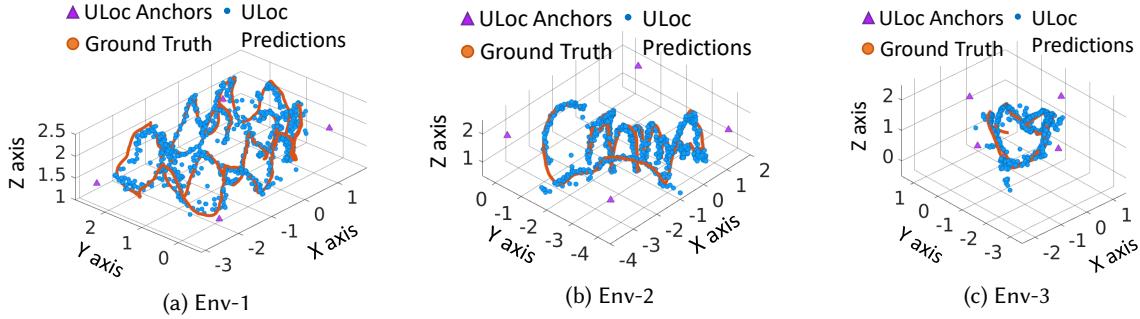


Fig. 8. Mobile tag tracking predictions: Shows the scatter plot of tracking predictions when the Tag is handheld and moved in (a) Env-1 (b) Env-2 (c) Env-3.

devices. Thus, in order to evaluate ULoc’s hardware design, we present Micro-benchmarks and Ablation studies in Section 4.6.

Test Environment: In the subsequent sections, we evaluate ULoc thoroughly in three real-world environments shown in Figure 7 and the tag moving at different velocities. Across these environments, we have collected over 10 days with a cumulative of over 300,000 locations. The first environment is a standard office setting with multiple natural reflectors spread out in the environment. Environment 2 and 3 are small and large warehouse-like indoor spaces with multiple strong reflectors as detailed in Figure 7(b, c). We deploy four ULoc Anchors in each of these environments. Each anchor is connected to a Raspberry Pi to record the UWB packets. Next, we set up the HTC Vive [22] system in the environment, and measure the positions of the anchors using the Vive trackers and laser markers. Next, we mount the ULoc tag to an HTC Vive [22] controller and use the locations furnished by the VR system as a viable source of ground truth [30] to compare ULoc’s 3D localization and tracking predictions. We also note that at anytime the testing happens within a $6\text{ m} \times 6\text{ m}$ space due to limitation of HTC vive, we move around HTC vive setup to test in large environment like Env-3, with proper groundtruth calibration.

Baselines: Within these environments, we compare ULoc with state-of-art and commercially available Pozyx [16] in both stationary and moving conditions. We also compare ULoc’s azimuth angle predictions with that of the PDoA kit’s [38]. Through this extensive data collection and testing, we have robustly tested ULoc’s 3D tracking and localization abilities. Apart from these end-to-end evaluations, we also study ULoc’s performance as we reduce the number of deployed anchors and reduce the number of antennas used on the ULoc Anchor. We also test the robustness of ULoc under different multipath scenarios and compare the 3D-localization performance when utilizing both FFT (Section 2) and MUSIC [97]. Finally, we discuss the power consumption of ULoc as compared to both SnapLoc [46] and two-way-ranging (TWR) based methods as outlined by Decawave [34].

4.1 Tracking Accuracy

Firstly, we evaluate ULoc’s tracking accuracy. Here, the tag is handheld and carried by a user in three different environments as shown in Figure 7 and at various velocities. The tracking results shown in Figure 9 and Figure 8 use 4 of ULoc’s Anchors for tracking the tag and 4 Pozyx Anchors to track Pozyx’s tag as well. From this CDF plot, we can see that Pozyx has $\{24, 79, 170\}$ cm at $\{50^{\text{th}}(\text{median}), 90^{\text{th}}, 99^{\text{th}}\}$ percentile tracking errors in Env1. While, ULoc achieves $\{8, 18, 44\}$ cm at $\{50^{\text{th}}(\text{median}), 90^{\text{th}}, 99^{\text{th}}\}$ percentile tracking errors in Env1, showing that ULoc outperforms state-of-the-art Pozyx system by $3\times$ overall in an office environment with a lot of natural reflectors including screens, cabinets, and desks.

Secondly, we also observe from Figure 9b we can see that in Env2, which is an open warehouse like scenario, that while Pozyx has a tracking error of $\{32, 100, 198\}$ cm at $\{50^{\text{th}}, 90^{\text{th}}, 99^{\text{th}}\}$ percentile. On the other hand, ULoc

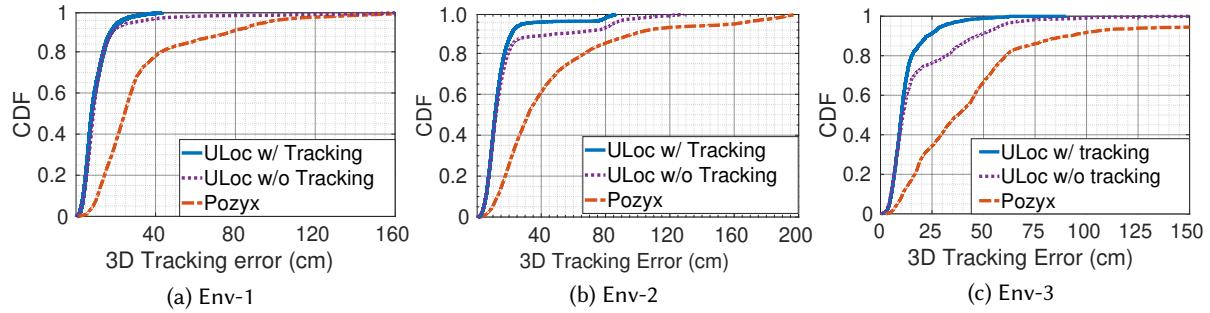


Fig. 9. (a) Mobile tag tracking errors: Shows the CDF of tracking errors when the Tag is handheld and moved in (a) Env-1 (b) Env-2 (c) Env-3. Here we compare ULoc against Pozyx [16] and also present the improvements provided by utilizing the unique tracking algorithm presented in Section 2.3

achieves $\{12, 22, 48\}$ cm at $\{50^{th}, 90^{th}, 99^{th}\}$ percentile tracking errors, showing up to $4\times$ better performance at the crucial 90^{th} percentile corner cases.

In addition, we test ULoc’s resilience to high multipath scenarios by evaluating it in Env3 that has a lot of reflectors in a much larger setup (20m \times 20m). We can see even in this scenario, where Pozyx has tracking errors of $\{38, 92, 430\}$ cm, ULoc achieves tracking errors of only $\{10, 23, 52\}$ cm at $\{50^{th}, 90^{th}, 99^{th}\}$ percentile respectively. Where in this scenario ULoc shows a performance improvement of $4\times$ for median and 90^{th} percentile, but outperforms Pozyx by $9\times$ at 99^{th} percentile. This shows the importance of ULoc’s tracking algorithm illustrated in section 2.4.

Finally, we can also see from Figure 9 that ULoc’s without tracking algorithm (FPI + 3D AoA based localization, does not take advantage of temporal tracking) achieves the same median performance as ULoc’s final tracking based algorithm. This shows the advantage of using FPI to identify the direct path for the majority of indoor scenarios. However, we see that the performance for ULoc’s with tracking outperforms ULoc’s without tracking performance at 90^{th} and 99^{th} percentiles by $2\times$ on average. Especially, we can see that in Env3 the rich multipath from the large environment significantly degrades both Pozyx’s and ULoc’s without tracking algorithms performance at 90^{th} and 99^{th} percentiles. While ULoc’s final tracking included algorithm improves ULoc’s without tracking performance by $3\times$ and Pozyx by $5\times$. Thus we can see that while ULoc’s final tracking algorithm makes it robust to adverse scenarios.

4.2 Accuracy under Different Velocities

Next, we analyze ULoc’s tracking performance under different mobility conditions, to show its robustness against varying velocities. Specifically, we showcase ULoc’s tracking accuracy under two velocity conditions – 0.7 m/s and 1.3 m/s, the velocity of delivery robots used commonly in warehouses [2]. From Figure 10(a), we see that the median ($90^{th}, 99^{th}$ percentile) errors for both these velocities are 12 cm (22 cm, 37 cm) and 14 cm (23 cm, 40 cm) respectively. To further enunciate the robustness of ULoc to different velocities of the tag, we plot the median tracking error binned into the nearest velocity of the tag for 8 different tag velocities and observed the results as shown in Figure 10(b). Here, we show the median tracking error vs velocity as an error-bar plot, with the error bars representing the standard deviation in tracking errors at the specific velocity bin. From this plot, we can see that median tracking accuracy does not go beyond 16cm for velocities as high as 2 m/sec which is well above the average moving velocities of assets or personnel in indoor scenarios, which is typically about 1 m/sec. This further demonstrates the real-time performance of ULoc, and the robustness of ULoc’s tracking algorithm in varied indoor settings.

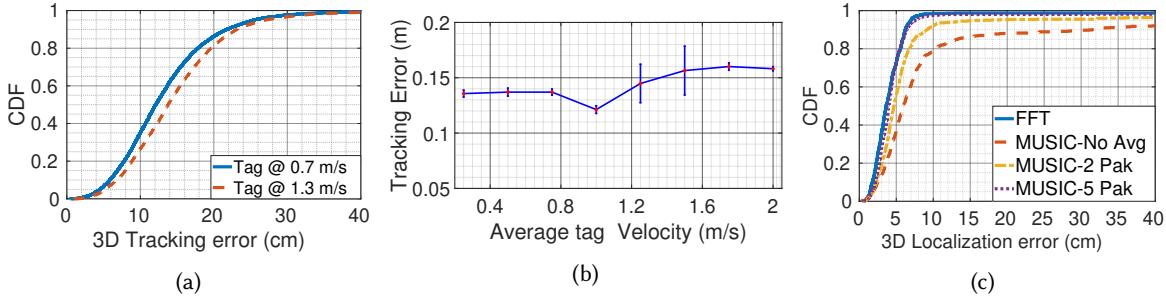


Fig. 10. (a) Tracking at different velocities Shows the tracking accuracy of tags moving at two different velocities, 0.7 m/s and 1.3 m/s. Note that there is very little change in tracking error with increased speeds. (b) Velocity vs Accuracy: Shows the tracking accuracy of ULoc when tag at different velocities (c) MUSIC vs. FFT: Shows the stationary-tag localization error when FFT and MUSIC is used. We see that while FFT does not require any packet averaging, MUSIC requires that the packets across the previous five time-instances be averaged.

Table 1. Comparison of Estimated Battery Life

	Method	Energy Per Location Est.	Est. Battery Life* for Tag @ 1Hz ($N_{AP} = 4$)	Est. Battery Life* for Tag @ 1kHz ($N_{AP} = 4$)	Tag Accuracy** (Stationary)
TWR Based [16]	$N_{AP} \times Rx + 2 \times Tx$	$N_{AP} \times 56 + 63 \mu J$	3.1 Months	2.3 Hours	6.8 cm (3D)
Concurrent Ranging [46]	$N_{AP} \times Rx$	$N_{AP} \times 56 \mu J$	4.0 Months	3.0 Hours	18 cm (2D)
TWR + AoA [91]	$2 \times Rx + 1 \times Tx$	$143 \mu J$	6.3 Months	4.6 Hours	6.5 cm (2D)
Concurrent AoA [49]	$N_{AP} \times Rx$	$N_{AP} \times 56 \mu J$	4.0 Months	3.0 Hours	30 cm (2D)
ULoc	1 × Tx	$31 \mu J$	2.4 Years	21.0 Hours	3.6 cm (3D)

*: For a CR2032 coin cell battery, assume no inactive state power consumption

**: Median accuracy reported in the corresponding work, under the setup that gives best accuracy

4.3 Localization Accuracy for Multiple Tags

Single Tag: First, we evaluate ULoc for stationary localization conditions to compare against many of the state of the art systems [46, 49, 69] which show similar static tag evaluation scenarios for a single UWB tag. We conduct these tests across 16 different stationary points with 4 of ULoc’s Anchors placed as shown in Figure 7(a) and Figure 11(a). We note that environment 1 is a standard office environment with multiple reflectors in the form of monitor screens and metal reflectors near the ceiling. ULoc’s localization accuracy in comparison with Pozyx is presented in Figure 11(b) where we achieve median, 90th, and 99th percentile accuracy of 3.6 cm, 6.3 cm and, 12.9 cm respectively as compared to Pozyx’s 6.8 cm, 16 cm and, 34.6 cm.

Multiple Tags: We also test for localization accuracy of multiple co-located tags in the environment. We implement a Time-division multiple access (TDMA) scheme to collect data from all the tags in static conditions. We collect data at 9 test-point in Env-1. In Figure 11(c), we show that there is little impact to performance when multiple tags are simultaneously deployed, with median and 90th percentile errors between 3.8 – 5.4 cm and 6.1 – 9.1 cm

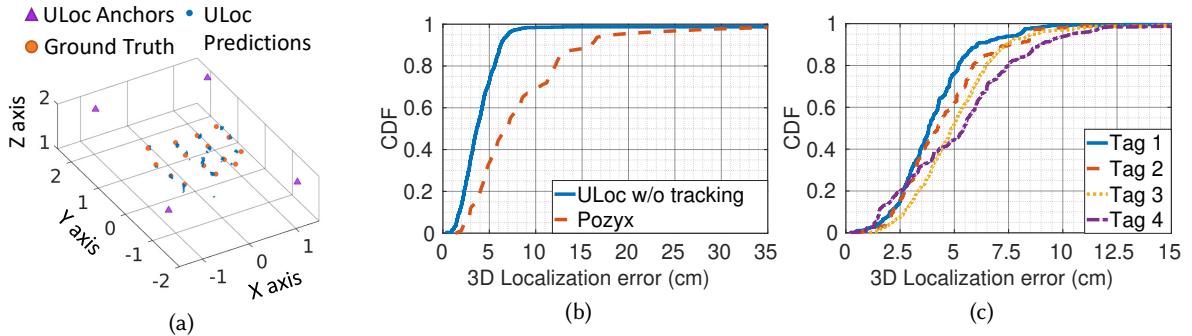


Fig. 11. (a, b) Stationary tag localization: Showcasing localization errors of a single stationary UWB Tag at the different locations in Env-1 in comparison with Pozyx tested under the same conditions. (c) Multi-tag localization: Shows the localization error CDF for four tags co-located in the same environment simultaneously. Note that increasing the number of tags in the environment has little impact on the localization accuracy.

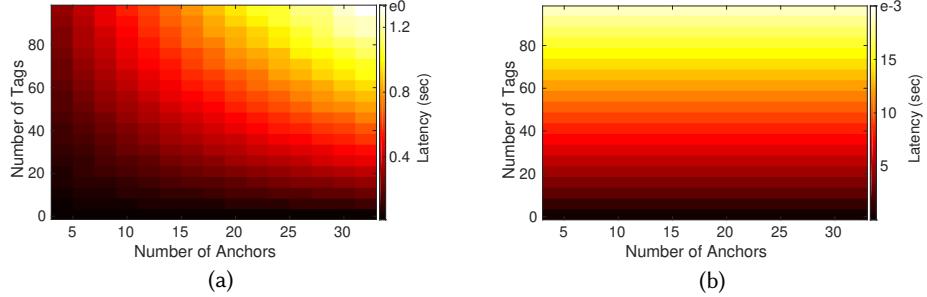


Fig. 12. Latency: Shows how latency increases with number of tags and anchors in a given environment for (a) TWR, where we can see that the latency is increasing exponentially going up to 1.4 sec (b) ULoc, where we can see that the latency is increasing linearly only with the number of tags in the environment and is completely independent of the number of anchors, thus reaching only a maximum latency of 20 msec.

MUSIC vs FFT: In Section 2, we make a design choice to use a ULoc’s FFT based algorithm to estimate the azimuth and polar angle of arrival using the CIR for the time-index corresponding to FPI. Another natural choice was MUSIC which was computationally expensive, but it starts working only after sufficient averaging of the packets [55]. Figure 10(c) shows the same behaviour for UWB, after averaging with 5 packets it performs close to the FFT based technique. This is primarily as the UWB has large bandwidth and therefore is less affected by multipath compared to low-bandwidth WiFi, and therefore even FFT without averaging is robust.

4.4 Latency Comparison

While we have seen that adding more tags to the system does not adversely affect its performance in terms of localization accuracy, we further evaluate ULoc latency efficacy with an increasing number of tags and anchors. To do this, we conduct a Poisson arrival model based MAC simulations and compare against standard TWR mechanism. The results of which are shown in Fig. 12. From this figure, we can see that while traditional TWR technologies’ latency scales up with respect to both number of Anchors and number of Tags, ULoc can do a simple time-division multiple access control to avoid scaling across increasing Anchors in the environment.

We can also see from Fig. 12 that while ULoc’s design has a maximum latency of 0.2secs for locating 100 tags across 32 different Anchors, a traditional TWR scales up to 1.5sec of latency. While these are the worst case scenario performances, we can see for a more general 4 Anchor setup, traditional TWR achieves 13 msec latency

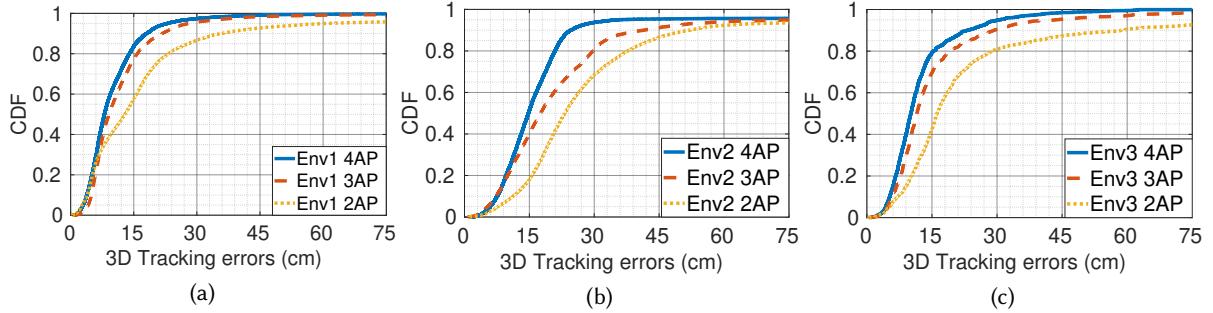


Fig. 13. Ablation: Varying number of Anchors: Shows the CDF of ULoc’s tracking error with increasing number of Anchors (AP) in (a) Env-1 (b) Env-2 (c) Env-3

for a single tag, and ULoc achieves a latency of just 1 msec. Though both of these values seem manageable, localization becomes delayed when you scale up to even just 100 tags, which would be a typical scenario in a warehouse tracking scenario. Where in the case of 100 tags to be localized, we can see that traditional TWR’s latency reaches 1.3 sec, ULoc still locates with a latency of 0.1 sec.

4.5 Power Consumption Comparison

Having robustly evaluated ULoc under a variety of conditions, next we turn to understanding ULoc’s unique design choice to provide low-power 3D localization via a single-packet transmission from the tag. In Table 1, we compare the power requirements and expected upper-limits on battery lives of the tag under different use cases. We use the power consumption measurements from the Decawave DW1000 datasheet [35] and the packet transmit and receive frame structure from the Decawave manual [6] to compute the total power consumed per location estimate by the Decawave DW1000 UWB chip. Specifically, we note that a single packet transmit requires $31\mu J$ and packet receive requires $56\mu J$. We combine these basic energy consumption computations with the localization protocols for ULoc, Concurrent Ranging-based [46], Concurrent AoA-based [49], hybrid of AoA and TWR [91], and TWR-based localization [16] to compute the energy consumed per location estimate at the tag. In ULoc, we use a single-packet transmission at the tag, in Concurrent Ranging-based and Concurrent AoA-based, the tag receives the number of packets equivalent to the number of Anchors, N_{AP} , in the environment and in TWR, there are 2 packet transmits and N_{AP} received packets. Note that for a hybrid TWR + AoA scheme, only a single Anchor is required for localization. Finally, we assume the nominal battery capacity (0.66 Wh) of the commonly used CR2032 lithium coin cell [39] to compute the best-case scenario of the battery-life of the tags under the different protocols. We summarize these protocols, the energy consumption, and estimated battery life in Table 1 at two different packet transmit frequencies. Specifically, most warehouse tracking would require a localization estimate at a frequency of 1 Hz, whereas high-mobility scenario would require a 1 kHz transmit frequency. Finally, we have also included a representative median localization accuracy as reported by the authors of the other papers. From the table, it is apparent ULoc provides superior accuracy while consuming the least amount of power for each location prediction

4.6 Micro-benchmarks and Ablation Studies

Now that we have seen the real-time, low-power, and cm-accurate 3D tracking performance of ULoc, we further evaluate the robustness of ULoc’s tracking with changing the number of Anchors deployed and robustness in ULoc’s 3D-AoA prediction with changing the number of antennas as compared to the Decawave PDoA kit[38].

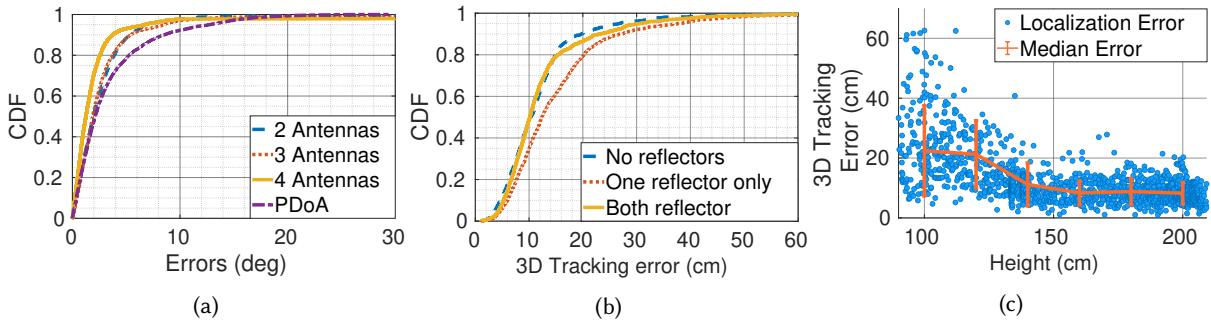


Fig. 14. (a) 3D-AoA performance: Shows the CDF of Angle errors for the Azimuth angle for ULoc and the PDoA Kit. We see that varying the number of antennas used for predicting the Azimuth angles degrades the performance marginally. Furthermore, we see that ULoc’s errors are smaller than the two-antenna PDoA Kit. (b) Robustness to reflectors: Shows the CDF of errors across three environment settings with different metal reflector setup in Env-3. ULoc has little change in performance with the addition of reflectors in the environment. (c) Shows the increase in error of as the tag is moved closer to the ground

Finally, we also study the change in ULoc’s tracking when additional reflectors are introduced in the environment.

Scalability in Anchors: One of the advantages of enabling 3D-AoA based 3D localization is that we can locate the user with as low as 2 anchors, and to see how well our algorithm fares with different numbers of anchors, we remove one anchor at a time and re-estimate the tag’s tracked location and estimate our tracking error as described in section 4.1. With this setup, we observe the two figures shown in Figure 13a,b for Env-1 and Env-2 respectively, wherein both the scenarios we observe that ULoc’s performance degrades significantly for a 2 Anchor system, but is more stable for a 3 Anchor and the usual 4 Anchor system for the given spaces. Where in Env-1, while a 4 Anchor deployment has a median tracking error of 8 cm, a 3 Anchor deployment has a median of 9 cm, but a 2 Anchor deployment has a median error of 13 cm. Similarly in Env-2 (Env-3), while a 4 Anchor deployment has a median tracking error of 12 cm (10.1 cm), the median tracking errors increase to 19 cm (11.5 cm) for 3 Anchors and 25 cm (16 cm) for 2 Anchors.

3D-AoA accuracy: In ULoc, we make an important contribution in tracking the 3D-AoA to improve the azimuth and polar angle predictions for a given point in space. We would like to further understand how ULoc’s 3D-AoA predictions compare with azimuth angle predictions of the PDoA kit [38]. For the PDoA kit testing, we mount the UWB tag, along with the VR controller, on a robot. We place the PDoA kit’s anchor at the same height as the UWB tag and move the robot in Env-2 (see Figure 7(b)). We collect the azimuth angle predictions from the PDoA kit and compare it against the ground truth provided by the co-located VR controller. Furthermore, in ULoc’s anchor, we reduce the number of antennas used in the horizontal and vertical antenna from 4 antennas to 2 antennas. We present the errors in the ULoc’s predicted 3D-AoA and PDoA kit’s predicted azimuth AoA in Figure 14(a). We find that reducing the number of antennas on ULoc’s Anchor reduced the median error (90th percentile) in the measured 3D-AoA from 1.4° (4.15°) to 2.1° (7.44°). Whereas, the 2-antenna PDoA kit achieves an accuracy of 2.32° (8.54°) at the median and (90th percentile)

Robustness to multipath: Finally, we would like to understand the performance of ULoc when different metal reflectors are added. In Section 4.1, we explored the tracking accuracy of ULoc in Env-3 (detailed in Figure 7(c, e)) when both the metal reflectors are present. To test ULoc’s performance in changing multipath conditions, we conducting tracking experiments with both metal reflector 1 and 2 removed, with only metal reflector 1 present,

and both metal reflectors present. In Figure 14(b) we present the results from these experiments. We see there is minimal difference in the tracking error when different reflectors are introduced.

5 RELATED WORK

ULoc builds the first low-latency, low-power, infrastructure-driven accurate 3D tracking and localization for UWB IoT devices. In this context, ULoc’s work overlaps with research in various fields ranging from RF-based localization and tracking to 3D visual tracking.

Light-based tracking: Camera-based [21, 104] or laser-based [22] systems which identify and track the different objects in the environment are unreliable when there are occlusions as is common-place in high-mobility conditions. This loss of tracking breaks away from R.1, which requires the infrastructure to be aware of the location of the tag in diverse conditions.

Acoustic-based tracking: To enable tracking in these visually-occluded conditions, there are many works which enable 3D acoustic-based tracking [42, 88] for VR applications. Unfortunately, acoustic-based tracking performance degrades with range. Moreover, in many industrial applications, ultrasonic noise in the frequency range which these acoustic tracking systems operate is prevalent [78]. This noise interference would degrade the 3D tracking performance of acoustic-based systems for industrial IoT applications, hence breaking away from R.2.

RF-based tracking: In contrast to visual and acoustic-based localization, RF-based localization can pass through most obstacles and is not corrupted by environmental noise. These advantages lead to a myriad of localization works with various RF protocols. While the research in WiFi-based localization [26, 44, 51, 55, 56, 59, 77, 80, 85, 86, 97–99] has been predominant in the last couple of decades due to its widespread deployment, these devices have limited bandwidth (20-160 MHz) compared to UWB (greater than 500 MHz). This relatively narrow bandwidth limits the localization accuracy to few decimeters. While the authors of WiCapture [56] achieve mm-accurate tracking (R.2) by taking advantage of a constant multipath present across multiple timestamps, this would not be possible in most of these indoor tracking scenarios and thus is not robust across various indoor environments (R.1). In the same vein, there have been works relying on BLE-based localization [27], but they suffer from high errors due to limited bandwidth as well. Hence, we find that both WiFi and BLE-based techniques are unable to meet R.1 to provide the sub-10 cm tracking accuracy.

Breaking away from active sensing and to enable a low-power tracking system, many systems have taken advantage of RFID-based backscatter [65, 89, 90, 101]. These works meet both the accuracy requirement (R.1) and provide a long operation lifetime for the tag (R.3). But, they have a very limited range and require the RFID reader to be within a few meters of the tag, hence requiring a high deployment density of RFID readers in the environment. This severely affects the scalability of these systems and fails to meet R.4. To combat the limited range of RFID-based systems, there have been other backscatter-based [69, 79] tracking approaches that require low power for operation and provide sub-10 cm accuracy. However, [69, 79] perform experiments in a more controlled setup with limited variance, compared to ULoc, where we evaluate it against ground truth measured from a VR tracker that is randomly moved by a human subject.

UWB Localization: To achieve the requirements R.1 to R.4, ULoc uses the Ultra-wideband (UWB) protocol which provides 1 nsec resolution in time owing to its wide bandwidth of 500 MHz. Taking advantage of this fine-time resolution, most of the existing UWB based localization algorithms utilize Two-Way Ranging (TWR) based localization [24, 29, 40, 41, 47, 48, 54, 57, 58, 60, 72, 105] which needs multiple packet exchanges to locate a single tag. These multi-packet exchange protocols significantly increase the latency to estimate a single location of a tag and hence fail to meet R.2. On the other hand, in ULoc we only need one ‘Blink’ from the tag to be transmitted that is simultaneously received by multiple anchors deployed in the environment. ULoc measures

Table 2. Comparison: Shows how ULoc is the only algorithm that can satisfy the increasing needs of 3D Indoor tracking amongst the existing work in UWB based indoor tracking. In the above power computation, 4 Anchors are assumed in systems which require multiple Anchors. Further details on power computation can be found in the appendix B.

	TWR Based [24, 40, 41, 60] [29, 47, 54]	Concurrent Ranging [31, 32, 46, 67, 87]	TWR + AoA [38, 91]	Concurrent AoA [49]	ULoc
3D Tracking	Yes	No	No	No	Yes
Accuracy	10s of cm	cm-accurate	10s of cm	10s of cm	cm-accurate
Tag Est. Battery Life	3.1 Months	4.0 Months	7.6 Months [38] 6.3 Months [91]	4.0 Months	2.4 Years
Infrastructure Driven	Yes	No	No	No	Yes
Latency	Very High	High	Medium	High	Low

the azimuth and polar angle of arrival of these incoming signals, collating this information across all the Anchors and localizes the tag in real-time.

On the other hand, recent systems are inspired by concurrent ranging and time-difference-of-arrival (TDoA) based localization algorithms [46, 61, 74, 84] that also perform real-time localization for UWB tags. However, these concurrent-ranging-based algorithms require the tags to receive and process multiple packets locally and localize themselves. This tag-driven localization is unfortunately power-hungry and compute-intensive on the tag's end as it is the tag that is receiving multiple transmissions and computing its location, thus does not meet R.3 and R.4. While these systems have deployment scenarios where the tag is accompanied by a computer like on a smartphone, they fail to scale up in an Industrial IoT scenario which ULoc targets. Furthermore, these systems are only deployed for 2D localization scenarios, thus departing from R.2 as well.

In contrast, TDoA schemes require only a single packet transmission at the tag and are infrastructure-driven. Unfortunately, the tedious synchronization required impedes TDoA systems from large-scale deployment and hence fails to meet R.4. Whereas, ULoc's anchor boards are connected to a common host or a central server which performs all the computations thus making the tag extremely low-power. Furthermore, ULoc does not suffer from the power/latency to accuracy trade-off since the 'Blink' from the tag is broadcast to all anchors in range, which then measure 3D-AoA and triangulate the tag. Thus, designing a multi-antenna anchor board and a 3D-AoA based tracking algorithm makes ULoc the first of its kind real-time and low power 3D tracking algorithm.

While, there have been few works [25, 43, 50, 62, 71, 83, 100] that have tried AoA based localization, most of these works have either been limited to simulations or VNA based experiments [43, 50, 62, 83, 100]. Recently, there have also been commercially available UWB-kits [15, 38], but these have been limited only to 2-D localization which does not satisfy R.2. Some works extend the TDoA approach by using statistical learning approaches to estimate AoA from multiple TDoA/TWR estimates [40, 41, 62], which is not robust to be tested across multiple different environments and so are not scalable (R.4).

Finally, there have been recent advances [25, 49, 94] that have tried estimating angles using Phase-difference-of-arrival (PDoA) and concurrent-AoA based techniques. Unfortunately, these works also perform the localization on the tag's end, and hence demand higher power (R.3) and compute for each location estimate (R.2). Further, all of these AoA based tracking systems are limited to 2D tracking or a 1D-AoA estimation, while with ULoc's unique board design we have enabled cm-accurate 3D tracking. While there have been a few 3D UWB localization algorithms [63, 70, 102] in the recent past, these systems have very high latency since they are based on ranging

based localization, and most of these works are in simulations with a few real-world experiments failing to meet R.2.

Thus we can see as summarized in Table 2 that while there exists a decade of research and development into UWB based localization and tracking, ULoc is one of its kind to achieve real-time, cm-accurate 3D tracking while providing long operational life for the tag.

6 LIMITATIONS AND FUTURE WORK

In this paper, we have proposed ULoc, a novel 3D tracking algorithm for UWB tags with contributions in both the hardware design towards the anchor board and the algorithm to achieve low-power, real-time, cm-grade 3D tracking. We have evaluated ULoc under different real-world settings under both static and moving conditions. While this work sparks one of its kind Low-power, real-time, and cm-accurate 3D tracking, there is further research in the following areas that can improve ULoc’s performance significantly:

- We note that despite the fine time resolution given by ULoc’s FPI, ground reflections at heights up to 1.3 meters can cause phase disruptions sufficient to interfere with ULoc’s 3D localization. Figure 14 (c) shows how the median and standard deviation of error at each height increases as we approach the ground. Initial thought as to the source of this error would be the non-uniform radiation pattern in the elevation 1 plane for the DWM1000 UWB chip [7]. Given this non-uniform single antenna gain, the full antenna array radiation pattern may cause unexpected destruction of the direct path for various elevations. With the direct path destroyed, prominent reflectors (such as ground reflections) may be picked up as the FPI, resulting in erroneous localization.
- The 2D antenna array design we have proposed is not optimal, there is research towards optimal antenna placements to achieve the highest field of view in terms of 3D localization and tracking.
- While in ULoc we achieve accurate tracking relying only on UWB estimates, similar to [16], we can integrate our solution with IMU measurements and develop much more accurate inertial measurements coupled tracking, leading towards VR tracking based applications.
- Additionally, the tag can send embedded data from the sensors attached to it (for example a temperature sensor or an Inertial Measurement Unit, IMU) within these *Blink* packets. These IMU measurements can be combined with the ULoc measurements to enable 6 DOF real-time tracking measurements.
- Further, tracking accuracy can also be improved by creating a rigid body of multiple UWB tags that can then be accurately tracked giving us 6 DOF measurements i.e., both location and orientation, as implemented by many of the existing vision-based VR tracking systems [22].

REFERENCES

- [1] 2021. AirFinder. <https://www.airfinder.com/>. (March 29 2021). Accessed: 2021-04-29.
- [2] 2021. Amazon Robotics warehouse robot datasheet. <https://robots.ieee.org/robots/kiva/>. (March 29 2021). Accessed: 2021-04-29.
- [3] 2021. Apple. <https://support.apple.com/en-in/guide/iphone/iph771fd0aad/ios>. (March 29 2021). Accessed: 2021-04-29.
- [4] 2021. Apple UWB-based Airtag. <https://www.apple.com/newsroom/2021/04/apple-introduces-airtag/>. (March 29 2021). Accessed: 2021-04-29.
- [5] 2021. Cisco. https://www.cisco.com/c/dam/global/sr_rs/assets/Wireless_in_retail.pdf. (March 29 2021). Accessed: 2021-04-29.
- [6] 2021. Decawave DW1000 User Manual. https://www.decawave.com/sites/default/files/resources/dw1000_user_manual_2.11.pdf. (March 29 2021). Accessed: 2021-04-29.
- [7] 2021. Decawave DWM1000 Module Datasheet. <https://www.decawave.com/sites/default/files/resources/DWM1000-Datasheet-V1.6.pdf>. (March 29 2021). Accessed: 2021-04-29.
- [8] 2021. Eliko. <https://www.eliko.ee/products/kio-rtls/>. (March 29 2021). Accessed: 2021-04-29.
- [9] 2021. Humatics. <https://timedomain.com/industries/research-and-education/scholar-ranging-localization/>. (March 29 2021). Accessed: 2021-04-29.
- [10] 2021. Infsoft. <https://www.infsoft.com/technology/positioning-technologies/ultra-wideband>. (March 29 2021). Accessed: 2021-04-29.
- [11] 2021. Impixon Indoor Asset tracking. <https://www.inpixon.com/technology>. (March 29 2021). Accessed: 2021-04-29.

- [12] 2021. Kinexion - Precise real time localization. <https://kinexion.com/>. (March 29 2021). Accessed: 2021-04-29.
- [13] 2021. Locatify. <https://locatify.com/>. (March 29 2021). Accessed: 2021-04-29.
- [14] 2021. MarvelMind. <https://marvelmind.com/>. (March 29 2021). Accessed: 2021-04-29.
- [15] 2021. NXP Ultra-wideband kit. <https://www.nxp.com/products/wireless/secure-ultra-wideband-uwb/trimension-uwb-development-kit-MK-UWB-DEV-KIT>. (March 29 2021). Accessed: 2021-04-29.
- [16] 2021. Pozyx. pozyx.io. (March 29 2021). Accessed: 2021-04-29.
- [17] 2021. Samsung Galaxy SmartTag+. <https://news.samsung.com/us/introducing-the-new-galaxy-smarttag-plus/>. (March 29 2021). Accessed: 2021-04-29.
- [18] 2021. Tenna Construction asset tracking. <https://www.tenna.com/use-cases/asset-tracking/>. (March 29 2021). Accessed: 2021-04-29.
- [19] 2021. Tile UWB tracker. <https://techcrunch.com/2021/01/05/tile-to-launch-to-launch-a-new-tracker-powered-by-ultra-wideband-technology/>. (March 29 2021). Accessed: 2021-04-29.
- [20] 2021. Ubisense. <https://ubisense.com/sensor-systems/>. (March 29 2021). Accessed: 2021-04-29.
- [21] 2021. Vicon. <https://www.vicon.com/>. (March 29 2021). Accessed: 2021-04-29.
- [22] 2021. Vive. <https://www.vive.com/us/product/#cosmos%20series>. (March 29 2021). Accessed: 2021-04-29.
- [23] Abraccon 2015. *Datasheet: ASTXR-12-38.400MHz-514054-T SMD TCXO*. Abraccon.
- [24] Abdulrahman Alarifi, AbdulMalik Al-Salman, Mansour Alsaleh, Ahmad Alnafessah, Suheer Al-Hadhami, Mai Al-Ammar, and Hend Al-Khalifa. 2016. Ultra wideband indoor positioning technologies: Analysis and recent advances. *Sensors* 16, 5 (2016), 707.
- [25] G Dickey Arndt, Phong H Ngo, Chau T Phan, Julia Gross, Jianjun Ni, and John Dusl. 2010. Ultra-Wideband Angle-of-Arrival Tracking Systems. (2010).
- [26] Roshan Ayyalasomayajula, Aditya Arun, Chenfeng Wu, Sanatan Sharma, Abhishek Rajkumar Sethi, Deepak Vasisht, and Dinesh Bharadia. 2020. Deep learning based wireless localization for indoor navigation. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–14.
- [27] Roshan Ayyalasomayajula, Deepak Vasisht, and Dinesh Bharadia. 2018. BLoc: CSI-based accurate localization for BLE tags. In *Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies*. 126–138.
- [28] Valentín Barral, Pedro Suárez-Casal, Carlos J Escudero, and José A García-Naya. 2019. Multi-sensor accurate forklift location and tracking simulation in industrial indoor environments. *Electronics* 8, 10 (2019), 1152.
- [29] Francisco Bonnin-Pascual and Alberto Ortiz. 2019. An UWB-based System for Localization inside Merchant Vessels. In *2019 24th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*. IEEE, 1559–1562.
- [30] Miguel Borges, Andrew Symington, Brian Coltin, Trey Smith, and Rodrigo Ventura. 2018. HTC vive: Analysis and accuracy improvement. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2610–2615.
- [31] Pablo Corbalán and Gian Pietro Picco. 2020. Ultra-wideband Concurrent Ranging. *ACM Transactions on Sensor Networks (TOSN)* 16, 4 (2020), 1–41.
- [32] Pablo Corbalán, Gian Pietro Picco, and Sameera Palipana. 2019. Chorus: UWB concurrent transmissions for GPS-like passive localization of countless targets. In *2019 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 133–144.
- [33] Decawave 2017. *TREK1000 software package*. Decawave. <https://www.decawave.com/software/>.
- [34] Decawave 2018. *The implementation of two-way ranging with the DW1000*. Decawave. Version 2.3.
- [35] Decawave 2020. *Datasheet: DW1000 IEEE802.15.4-2011 UWB Transceiver*. Decawave. Version 2.22.
- [36] Diodes Incorporated 2011. *Datasheet: AP1117 1A Low Dropout Regulator*. Diodes Incorporated. Document number: DS31009 Rev. 21 - 2.
- [37] Diodes Incorporated 2017. *Datasheet: AP2112 600mA CMOS LDO REGULATOR WITH ENABLE*. Diodes Incorporated. Document number: DS39724 Rev. 2 - 2.
- [38] Igor Dotlic, Andrew Connell, Hang Ma, Jeff Clancy, and Michael McLaughlin. 2017. Angle of arrival estimation using decawave DW1000 integrated circuits. In *2017 14th Workshop on Positioning, Navigation and Communications (WPNC)*. IEEE, 1–6.
- [39] Energizer [n. d.]. *Datasheet: CR2032 Lithium Coin cell*. Energizer. Doc. 2032NA0618.
- [40] Enrique García, Pablo Poudereux, Álvaro Hernández, Juan Jesús García, and Jesús Ureña. 2013. DS-UWB indoor positioning system implementation based on FPGAs. *Sensors and Actuators A: Physical* 201 (2013), 172–181.
- [41] Enrique García, Pablo Poudereux, Álvaro Hernández, Jesús Ureña, and David Gualda. 2015. A robust UWB indoor positioning system for highly complex environments. In *2015 IEEE International Conference on Industrial Technology (ICIT)*. IEEE, 3386–3391.
- [42] Linfei Ge, Qian Zhang, Jin Zhang, and Qianyi Huang. 2020. Acoustic Strength-based Motion Tracking. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 4 (2020), 1–19.
- [43] Waldemar Gerok, Mohamed El-Hadidy, Sondos Alaa El Din, and Thomas Kaiser. 2010. Influence of the real UWB antennas on the AoA estimation based on the TDoA localization technique. In *IEEE Middle East Conference on Antennas and Propagation (MECAP 2010)*. IEEE, 1–6.
- [44] Jon Gjengset, Jie Xiong, Graeme McPhillips, and Kyle Jamieson. 2014. Phaser: Enabling Phased Array Signal Processing on Commodity Wi-Fi Access Points. *MobiCom* (2014).

[45] Vineet Gokhale, Gerardo Moyers Barrera, and R Venkatesha Prasad. 2021. FEEL: Fast, Energy-Efficient Localization for Autonomous Indoor Vehicles. *arXiv preprint arXiv:2102.00702* (2021).

[46] Bernhard Großwindhager, Michael Stocker, Michael Rath, Carlo Alberto Boano, and Kay Römer. 2019. SnapLoc: An ultra-fast UWB-based indoor localization system for an unlimited number of tags. In *2019 18th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 61–72.

[47] Bernhard Großwindhager, Michael Rath, Josef Kulmer, Mustafa S Bakr, Carlo Alberto Boano, Klaus Witrisal, and Kay Römer. 2018. SALMA: UWB-based single-anchor localization system using multipath assistance. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*. 132–144.

[48] Ismail Guvenc, Chia-Chin Chong, and Fujio Watanabe. 2007. NLOS identification and mitigation for UWB localization systems. In *2007 IEEE Wireless Communications and Networking Conference*. IEEE, 1571–1576.

[49] Milad Heydariaan, Hossein Dabirian, and Omprakash Gnawali. 2020. AnguLoc: Concurrent angle of arrival estimation for indoor localization with UWB radios. In *2020 16th International Conference on Distributed Computing in Sensor Systems (DCOSS)*. IEEE, 112–119.

[50] Naohiko Iwakiri and Takehiko Kobayashi. 2008. Ultra-wideband time-of-arrival and angle-of-arrival estimation using transformation between frequency and time domain signals. *Journal of Communications* 3, 1 (2008), 12–19.

[51] Kiran Joshi, Steven Hong, and Sachin Katti. 2013. PinPoint: Localizing Interfering Radios (NSDI).

[52] Mathias Källström, Sondre Berdal, and Suhas Govind Joshi. 2015. Designing an indoor navigation system for elderly people's capabilities. In *International Conference on Human Aspects of IT for the Aged Population*. Springer, 435–445.

[53] Eirini Karapistoli, Fotini-Niovi Pavlidou, Ioannis Gragopoulos, and Ioannis Tsetsinas. 2010. An overview of the IEEE 802.15. 4a standard. *IEEE Communications Magazine* 48, 1 (2010), 47–53.

[54] Benjamin Kempke, Pat Pannuto, and Prabal Dutta. 2016. Harmonium: Asymmetric, bandstitched UWB for fast, accurate, and robust indoor localization. In *2016 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 1–12.

[55] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. SpotFi: Decimeter Level Localization Using Wi-Fi (SIGCOMM).

[56] Manikanta Kotaru and Sachin Katti. 2017. Position tracking for virtual reality using commodity WiFi. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 68–78.

[57] Sivanand Krishnan, Pankaj Sharma, Zhang Guoping, and Ong Hwee Woon. 2007. A UWB based localization system for indoor robot navigation. In *2007 IEEE International Conference on Ultra-Wideband*. IEEE, 77–82.

[58] Michael Kuhn, Cemin Zhang, Brandon Merkl, Depeng Yang, Yazhou Wang, Mohamed Mahfouz, and Aly Fathy. 2008. High accuracy UWB localization in dense indoor environments. In *2008 IEEE International Conference on Ultra-Wideband*, Vol. 2. IEEE, 129–132.

[59] Swarun Kumar, Stephanie Gil, Dina Katabi, and Daniela Rus. 2014. Accurate Indoor Localization with Zero Start-up Cost (*MobiCom*).

[60] Anton Lederergerber and Raffaello D'Andrea. 2019. Ultra-Wideband Angle of Arrival Estimation Based on Angle-Dependent Antenna Transfer Function. *Sensors* 19, 20 (2019), 4466.

[61] Anton Lederergerber, Michael Hamer, and Raffaello D'Andrea. 2015. A robot self-localization system using one-way ultra-wideband communication. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3131–3137.

[62] Yong Up Lee. 2011. Weighted-average based aoa parameter estimations for LR-UWB wireless positioning system. *IEICE transactions on communications* 94, 12 (2011), 3599–3602.

[63] Jiaxin Li, Yingcai Bi, Kun Li, Kangli Wang, Feng Lin, and Ben M Chen. 2018. Accurate 3d localization for mav swarms by uwb and imu fusion. In *2018 IEEE 14th International Conference on Control and Automation (ICCA)*. IEEE, 100–105.

[64] Yanjia Luo and Choi Look Law. 2009. Angle-of-arrival estimation with array in a Line-of-Sight indoor UWB-IR. In *2009 7th International Conference on Information, Communications and Signal Processing (ICICS)*. IEEE, 1–5.

[65] Yunfei Ma, Nicholas Selby, and Fadel Adib. [n. d.]. Drone Relays for Battery-Free Networks. In *SIGCOMM*.

[66] Rainer Mautz. 2012. Indoor positioning technologies. (2012).

[67] Michael McLaughlin and Jaroslaw Niewczas. 2020. Ranging with simultaneous frames. (Nov. 5 2020). US Patent App. 16/761,160.

[68] Microchip 2020. *Datasheet: Microchip PL133-37 Low-Power, 1.62V to 3.63V, 1 MHz to 150 MHz, Inverting 1:3 Fanout Buffer IC*. Microchip. Revision A.

[69] Rajalakshmi Nandakumar, Vikram Iyer, and Shyamnath Gollakota. 2018. 3d localization for sub-centimeter sized devices. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*. 108–119.

[70] Dongchen Ni, Octavian Adrian Postolache, Chao Mi, Meisu Zhong, and Yongshuang Wang. 2019. UWB indoor positioning application based on Kalman filter and 3-D TOA localization algorithm. In *2019 11th International Symposium on Advanced Topics in Electrical Engineering (ATEE)*. IEEE, 1–6.

[71] Idnin Pasya, Naohiko Iwakiri, and Takehiko Kobayashi. 2014. Joint direction-of-departure and direction-of-arrival estimation in a UWB MIMO radar detecting targets with fluctuating radar cross sections. *International Journal of Antennas and Propagation* 2014 (2014).

[72] A. Poulose, Z. Emersic, O. Steven Eyobu, and D. Seog Han. 2020. An Accurate Indoor User Position Estimator For Multiple Anchor UWB Localization. In *2020 International Conference on Information and Communication Technology Convergence (ICTC)*. 478–482. <https://doi.org/10.1109/ICTC49870.2020.9289338>

[73] Raspberry Pi (Trading) Ltd 2019. *Datasheet: Raspberry Pi Compute Module 3+*. Raspberry Pi (Trading) Ltd. Release 1.

[74] Nathan C Rowe, Aly E Fathy, Michael J Kuhn, and Mohamed R Mahfouz. 2013. A UWB transmit-only based scheme for multi-tag support in a millimeter accuracy localization system. In *2013 IEEE Topical Conference on Wireless Sensors and Sensor Networks (WiSNet)*. IEEE, 7–9.

[75] R. Roy and T. Kailath. 1989. ESPRIT-estimation of signal parameters via rotational invariance techniques. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 37, 7 (1989), 984–995. <https://doi.org/10.1109/29.32276>

[76] Ralph Schmidt. 1986. Multiple emitter location and signal parameter estimation. *IEEE transactions on antennas and propagation* 34, 3 (1986), 276–280.

[77] Souvik Sen, Božidar Radunovic, Romit Roy Choudhury, and Tom Minka. 2012. You are facing the Mona Lisa: Spot localization using PHY layer information. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*. ACM, 183–196.

[78] Bożena Smagowska. 2013. Ultrasonic noise sources in a work environment. *Archives of Acoustics* 38, 2 (2013), 169–176.

[79] Elahe Soltanaghaei, Adwait Dongare, Akarsh Prabhakara, Swarun Kumar, Anthony Rowe, and Kamin Whitehouse. 2021. TagFi: Locating Ultra-Low Power WiFi Tags Using Unmodified WiFi Infrastructure. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 1 (2021), 1–29.

[80] Elahe Soltanaghaei, Avinash Kalyanaraman, and Kamin Whitehouse. 2018. Multipath Triangulation: Decimeter-level WiFi Localization and Orientation with a Single Unaided Receiver. In *MobiSys*.

[81] STMicroelectronics 2019. *Datasheet DS9826: STM32F072xB ARM®-based 32-bit MCU*. STMicroelectronics. Revision 6.

[82] STMicroelectronics 2021. *Datasheet DS10693: STM32F446xC/E Arm Cortex-M4 32-bit MCU*. STMicroelectronics. Revision 10.

[83] Ananth Subramanian. 2007. UWB linear quadratic frequency domain frequency invariant beamforming and angle of arrival estimation. In *2007 IEEE 65th Vehicular Technology Conference-VTC2007-Spring*. IEEE, 614–618.

[84] Janis Tiemann, Fabian Eckermann, and Christian Wietfeld. 2016. Atlas—an open-source tdoa-based ultra-wideband localization system. In *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 1–6.

[85] Michaela C Vanderveen, Constantinos B Papadias, and Arogyaswami Paulraj. 1997. Joint angle and delay estimation (JADE) for multipath signals arriving at an antenna array. *IEEE Communications letters* 1, 1 (1997), 12–14.

[86] Deepak Vasisht, Swarun Kumar, and Dina Katabi. 2016. Decimeter-Level Localization with a Single Wi-Fi Access Point (NSDI).

[87] Davide Vecchia, Pablo Corbalán, Timofei Istomin, and Gian Pietro Picco. 2019. Playing with fire: Exploring concurrent transmissions in ultra-wideband radios. In *2019 16th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. IEEE, 1–9.

[88] Anran Wang and Shyamnath Gollakota. 2019. Millisonic: Pushing the limits of acoustic motion tracking. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–11.

[89] Jue Wang and Dina Katabi. 2013. Dude, Where's My Card?: RFID Positioning That Works with Multipath and Non-line of Sight (SIGCOMM).

[90] Jue Wang, Deepak Vasisht, and Dina Katabi. 2014. RF-IDraw: virtual touch screen in the air using RF signals. *ACM SIGCOMM Computer Communication Review* 44, 4 (2014), 235–246.

[91] Tianyu Wang, Hanying Zhao, and Yuan Shen. 2020. An efficient single-anchor localization method using ultra-wide bandwidth systems. *Applied Sciences* 10, 1 (2020), 57.

[92] Zhihua Wang, Zhaochu Yang, and Tao Dong. 2017. A review of wearable technologies for elderly care that can accurately track indoor position, recognize physical activities and monitor vital signs in real time. *Sensors* 17, 2 (2017), 341.

[93] Greg Welch and Gary Bishop. 1995. *An introduction to the kalman filter*. Chapel Hill, NC. Technical Report. USA, Tech. Rep.

[94] Stijn Wielandt and Lieven De Strycker. 2017. Indoor multipath assisted angle of arrival localization. *Sensors* 17, 11 (2017), 2522.

[95] Zhiqiang Xiao and Yong Zeng. 2020. An overview on integrated localization and communication towards 6G. *arXiv preprint arXiv:2006.01535* (2020).

[96] Yaxiong Xie, Jie Xiong, Mo Li, and Kyle Jamieson. 2018. mD-Track: Leveraging Multi-Dimensionality in Passive Indoor Wi-Fi Tracking. *arXiv preprint arXiv:1812.03103* (2018).

[97] Jie Xiong and Kyle Jamieson. 2013. ArrayTrack: A Fine-grained Indoor Location System (NSDI).

[98] Jie Xiong, Kyle Jamieson, and Karthikeyan Sundaresan. 2014. Synchronicity: Pushing the Envelope of Fine-grained Localization with Distributed Mimo. In *HotWireless*.

[99] Jie Xiong, Karthikeyan Sundaresan, and Kyle Jamieson. 2015. ToneTrack: Leveraging Frequency-Agile Radios for Time-Based Indoor Wireless Localization (*MobiCom*).

[100] Jun Xu, Maode Ma, and Choi Look Law. 2008. AOA cooperative position localization. In *IEEE GLOBECOM 2008-2008 IEEE Global Telecommunications Conference*. IEEE, 1–5.

[101] Lei Yang, Yekui Chen, Xiang-Yang Li, Chaowei Xiao, Mo Li, and Yunhao Liu. 2014. Tagoram: Real-time Tracking of Mobile RFID Tags to High Precision Using COTS Devices (*MobiCom*).

[102] Yonggang Zhang, Feng Shen, Qi Liu, and Wenqiang Li. 2020. A UWB 3D Localization Algorithm Based on Residual Weighting. In *2020 IEEE International Conference on Mechatronics and Automation (ICMA)*. IEEE, 309–313.

- [103] Zijian Zhang, Hanying Zhao, and Yuan Shen. 2019. High-efficient ranging algorithms for wireless sensor network. In *2019 11th International Conference on Wireless Communications and Signal Processing (WCSP)*. IEEE, 1–6.
- [104] Haidi Zhu, Haoran Wei, Baoqing Li, Xiaobing Yuan, and Nasser Kehtarnavaz. 2020. A Review of Video Object Detection: Datasets, Metrics and Methods. *Applied Sciences* 10, 21 (2020), 7834.
- [105] Lukasz Zwirello, Tom Schipper, Marlene Harter, and Thomas Zwick. 2012. UWB localization system for indoor applications: Concept, realization and analysis. *Journal of Electrical and Computer Engineering* 2012 (2012).

A FPI + 2D FFT BASED 3D AOA ESTIMATION

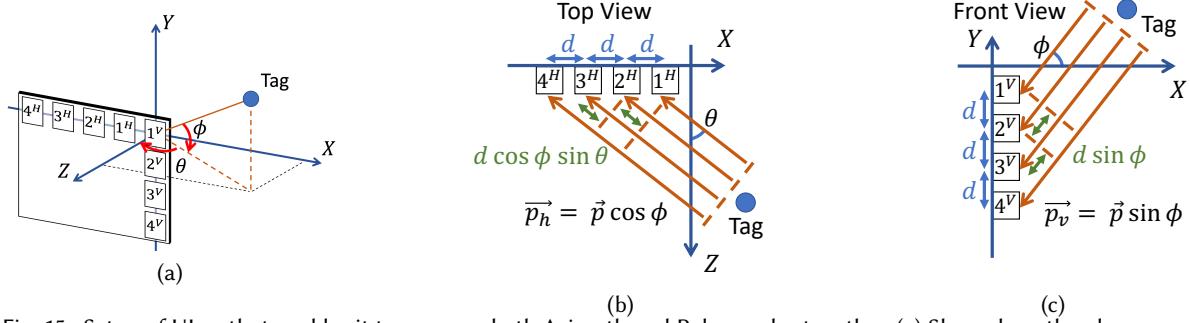


Fig. 15. Setup of ULoc that enables it to measure both Azimuth and Polar angles together. (a) Shows how the planar wave incident on the antenna array can be projected onto two perpendicular planes. (b) Shows zoomed-in version on how the signal travels an additional $d \cos \phi \sin \theta$ for each additional antenna on the horizontal array, which can be extended to see that (c) the signal travels an additional distance of $d \sin \phi$ distance for each additional antenna along with the vertical array.

To define our algorithm more formally, let us consider the scenario shown in Figure 15. Specifically, under the far-field assumption the ideal channel impulse response (CIR), where n is the time-index, $h_i^v(n)$, for the i^{th} antenna ($\forall i = 1, 2 \dots, N_{\text{ant}-v}$) on the vertical antenna array, and $h_j^h(n)$, for the j^{th} antenna ($\forall j = 1, 2 \dots, N_{\text{ant}-h}$) on the horizontal antennas array, can be shown [27, 55, 97] to be

$$h_i^v(n) = a_i \exp \left(-i2\pi f_c \frac{r + (i-1)d \sin \phi}{c} \right), \quad h_j^h(n) = a_j \exp \left(-i2\pi f_c \frac{r + (j-1)d \sin \theta \cos \phi}{c} \right) \quad (4)$$

for a tag located at distance r from the anchor board transmitted at central frequency f_c . Where c is the speed of light, $i = \sqrt{-1}$, and a_i, a_j are the corresponding attenuation constants across the antennas. For this given channel impulse response measurements on a given tap n , one can define a simple transform based on Eq 4 to get Eq 5. The intuition in defining the transform is to balance the phase term in the exponent, for example, $\phi_l = \phi$ at index n , all the terms of the first summation would add in-phase leading to $N_{\text{ant}-v} \cdot a_i$, otherwise it would be destructive.

$$P(\theta_k, \phi_l, n) = \left| \sum_{i=1}^{N_{\text{ant}-v}} h_i^v(n) \exp \left(i2\pi \frac{(i-1)d \sin \phi_l}{c} \right) + \sum_{j=1}^{N_{\text{ant}-h}} h_j^h(n) \exp \left(i2\pi \frac{(j-1)d \sin \theta_k \cos \phi_l}{c} \right) \right| \quad (5)$$

where $P(\theta_k, \phi_l, n)$ is the likelihood at the direction of arrival of (θ_k, ϕ_l) corresponding to the time-index n , of the measured CIR. The above transform is simply a Fast Fourier transform. Furthermore, due to the design of the antenna arrays, under no multi-path, there is always a unique solution for the actual direct path's 3D-AoA, $(\hat{\theta}, \hat{\phi})$ as

$$(\theta^{\text{UWB}}, \phi^{\text{UWB}}) = \underset{-\pi/2 \leq \theta \leq \pi/2}{\underset{-\pi/2 \leq \phi \leq \pi/2}{\text{argmax}}} P(\theta, \phi) \quad (6)$$

This is under the assumption that $\theta \in [-\pi/2, \pi/2]$ and $\phi \in [-\pi/2, \pi/2]$, no carrier frequency offset, and $d \leq \frac{\lambda}{2}$, where λ is the wavelength of the carrier frequency. For any antenna separation $d > \frac{\lambda}{2}$ there would be aliasing, which is tackled in the hardware design of ULoc's anchor which is further described in section 3.1.

B POWER CONSUMPTION

Table 3. Energy Consumption at Each State

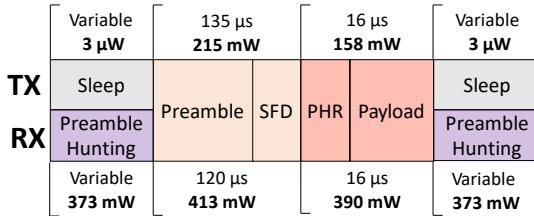


Fig. 16. Tx/Rx States and Power Consumption

	UWB State	Current*	Duration	Energy
Active Tx	Preamble + SFD	65 mA	135 μ s	28.96 μ J
	PHR + Payload (12 byte frame)	48 mA	16 μ s	2.54 μ J
Inactive Tx	Sleep	1 μ A	Variable	-
Active Rx	Preamble + SFD	125 mA	120 μ s	49.5 μ J
	PHR + Payload (12 byte frame)	118 mA	16 μ s	6.23 μ J
Inactive Rx	100% Preamble Hunt	113 mA	Variable	-
	50% Preamble Sniff	62.5 mA	Variable	-

*: Current measured under 3.3V

In Table 1, we see that in different related works, the amount of Tx and Rx required for each location estimate is different, thus yielding different power consumption. In this section, we provide details on how the power consumption and battery life are estimated.

According to IEEE 802.15.4 standard, each UWB packet starts with a preamble, then followed by the start of frame delimiter (SFD), PHY Header (PHR), and payload [53]. As shown in Fig. 16, on the Tx side, the UWB radio wakes up from sleep, transmits the four segments in order, and goes back to sleep again. On the Rx side, besides receiving each segment, a preamble hunting stage is required to identify the existence of the packet. We denote the transmitting and receiving the actual packet as *active* Tx/Rx, and the rest as *inactive* Tx/Rx.

The power consumption on each of the states, along with their duration, are acquired from Decawave DW1000 Datasheet [35]. The Tx/Rx profile is measured with 16MHz Pulse Repetition Frequency (PRF), data rate 6.8Mb/s, Channel 2, Preamble length 128 symbols, and 12-byte frame for header + payload. Note that the inactive states do not have a fixed duration – they are dependent on the localization rate and the MAC protocol.

Active Tx / Rx From Fig. 16 we see that transmitting or receiving the four segments within a UWB packet has a fixed current consumption and duration, thus we can individually calculate the energy consumed in Tx / Rx for each of the four-segment, and sum up to obtain the total energy consumed for one active Tx / Rx (Table 3).

Inactive Tx / Rx For transmitting, the inactive state is usually *sleep* between active states, whereas for receiving, besides *sleep*, the inactive state also includes a high power-consumption *preamble hunt* which continuously Rx and search for the existence of preamble that denotes the start of a UWB packet.

The duration that the UWB radio stays in inactive states is defined by the localization rate and MAC protocol and is highly variable between different related works. Thus we do not consider the inactive state power consumption in our energy per location estimate calculation. However, we note that in ULoc, as the tag does not perform Rx, the inactive power consumption is very low as compared to other related work that requires Rx.

Estimate Battery Life The battery life is estimated on a CR2032 0.66Wh coin-cell battery [39], which converts to 2376J of total energy.

C DATA EXTRACTION AND PROCESSING PIPELINE

C.1 Extracting Data:

As discussed in section 2.2, ULoc anchor uses channel impulse response (CIR) at the rising edge of the direct path signal coming from ULoc tag to estimate 3-D AoA and localize the tag. When a UWB packet arrives at a DW1000 transceiver, it holds the accumulated CIR data in a bank of its memory, while Decawave's proprietary leading-edge detection (LDE) algorithm reports an index in the stored CIR array to the "leading edge" defining the first receiving ray. The CIR is reported in 1 ns resolution, whereas the first path index (FPI) reported by the LDE algorithm is a fractional number that has a 15.65 ps resolution, 64 times more than the CIR taps [6]. After that, the transceiver asserts its IRQ to notify the MCU of the Rx event. Upon notification, the MCU obtains the required data from the respective transceiver.

To obtain the CIR around FPI, the MCU first reads the transceiver's register *FP_INDEX* via SPI, obtains FPI, and round it to the nearest integer. Next, it accesses the accumulator memory *ACC_MEM* from an offset of *FPI* – 1 to *FPI* + 5, for a total of 7 CIR taps (roughly 7 ns of channel information around the first path). In addition, the CIR extracted on-board at each transceiver has an additional delay corresponding to the accurate identification of the start-of-frame delimiter (SFD) [6], that is reported by the DW1000 chip. This SFD's phase can be used to compensate out the phase offset between multiple DW1000 transceivers caused by this delay when subtracted from the CIR. We access the SFD phase information from register *RCPHASE*.

Furthermore, we obtain DW1000 temperature information from register *SAR_LTEMP* for phase difference calibration as detailed in section 3.4. We also obtain the tag's unique ID from Rx frame data buffer register *RX_BUFFER*, so that we know which tag this data is tied to.

Output Data to Infrastructure:

The anchor MCU has all necessary data extracted from each of the 8 transceivers. To estimate location, these data need to be gathered together from multiple anchors. The anchor firmware includes a USB communications device class (CDC) driver that allows the MCU to appear as an emulated serial port and send data upon connecting to a host computer.

A straightforward way to obtain data is to connect each anchor to a central host computer via USB. However, this solution will quickly become infeasible as the number of anchors or their distances goes up. Not only will the data rate be limited by the flooded I/O on the central computer, but also the long USB cables and a limited amount of USB ports on the computer become issues to the infrastructure. To make the system scalable, we attach dedicated Raspberry Pi single-board computers [73] to each of the anchors, so that data from each anchor is received and computed separately at the edge. Each Raspberry Pi reads the data from an anchor, performs all corrections needed, and estimates 3-D AoA. We connect all the Raspberry Pis to the same Wi-Fi network so they can transmit the 3-D AoA information to the infrastructure and the tag's location is computed by building the algorithm on, to be solved.

C.2 Data Processing Pipeline

As discussed in section 3.2, we have acquired from each anchor the raw data required to perform localization, including 7 taps of CIR with 1 ns resolution, the FPI with 15.65 ps resolution, and SFD phase for compensation. We also assume wireless calibration as discussed in section 3.4 has been performed and a zero AoA phase difference calibration data exist. In this sub-section, we list down the data processing steps required for the infrastructure to get accurate, reliable, real-time localization for the ULoc tags.

1. CIR Cleaning and Upsampling: The packets belonging to different tags are first filtered out by tag ID, removing the packets which do not belong to our data case. The CIR is upsampled by 8 times using an 8-tap raised-cosine filter, to take advantage of the 64 times more resolution FPI has. We have chosen to upsample 8 times

because we have observed that upsampling 64 times gives a negligible performance boost while significantly increase the amount of processing time.

2. CIR Compensation: As discussed in section 3.2, the SFD phase can be used to compensate out the phase offset between multiple transceivers. Thus for each transceiver, we subtract its SFD phase from its CIR taps. Then, we further subtract the zero AoA phase offset acquired in wireless calibration. Now the CIR is ready for 3D-AoA estimation.

3. 3D-AoA Estimation: First, as discussed in section 2.2, FFT is applied on the CIR across antenna at FPI tap, to obtain the AoA likelihood profile for each anchor. Then, as seen in section 2.3, the likelihood profile is processed by the Temporal Tracker and Close Peak Finder to obtain two 3D-AoA estimations. A weighted average is applied to these estimations to obtain a final ULoc 3D-AoA.

4. 3D Localization: With the cleaned 3D-AoA's obtained from the last step on multiple ULoc anchor, triangulation is used to solve for the location of the tag (section 2.4).